

Laws and Skills – An Inferential Diagnosis and Defense

Inauguraldissertation
zur Erlangung des Grades eines Doktors der Philosophie
im Fachbereich 8
der Johann-Wolfgang-Goethe-Universität
zu Frankfurt am Main

vorgelegt von
Hong Zhou
aus: Xi'an, China

Eingereicht am 10.11.2010

1. Gutachter: Prof. Dr. T. E. Zimmermann
2. Gutachter: Prof. Dr. P. Stekeler-Weithofer
3. Gutachter: Prof. Dr. A. Fuhrmann

Table of Contents

Prologue	3
Part I: Semantic Preludes.....	8
§1 Possible Worlds and Semantic Modeling.....	9
§1.1 Semantic Modeling vs. Semantic Explanation	10
§1.2 Possible Worlds I – Intensional Semantic Modeling.....	11
§1.3 Conceptions of Possible Worlds	22
§1.4 Possible Worlds II – Modal Vocabulary.....	36
§1.5 Conclusions.....	43
§2 Semantic Functionalism – A Selective Overview.....	45
§2.1 Introduction: Two Sources and Four Choices	46
§2.2 Two Sources – Dummett and Sellars.....	48
§2.3 Content of Thoughts vs. Meaning of Expressions	80
§2.4 Which Semantic Functionalism?	85
§2.5 Conclusions.....	104
Part II Semantic Functionalism & The Problems of Laws	106
Introduction to Part II.....	107
§3 Inferentialists on Nomological Necessities & Explicit Making	108
§3.1 Introduction	109
§3.2 Dummett - Anti-Realism about Abilities.....	110
§3.3 Sellars and Brandom on Material Rules of Inference.....	112
§3.4 Brandom’s Expressivism	135
§4 Harmony – A Diagnosis	154
§4.1 The Idea Of Harmony	155
§4.2 The Problem of Laws of Nature	165
§4.3 Realism vs. Anti-Realism About Laws of Nature	174
Part III: Laws and Skills – A Therapy.....	192
§5 Elements of Nomological Practice: Experiments, Practical Deliberations, and Skills.....	193
§5.1 Introduction	194
§5.2 Desiderative Inferences and the Practical Deliberation Principle.....	199
§5.3 The Skill-by-Experiment Principle	218
§5.4 The Problems of Disharmony Revisited.....	236
§6 Objections.....	242
§6.1 Introduction	243
§6.2 Perception, Empiricism, and Transcendental Arguments.....	244
§6.3 Non-Perceptual Elements.....	252
§6.4 Skepticisms.....	276
Part IV: Beyond Therapy	278
§7 Making Explicit & Epistemé.....	279
§7.1 Introduction	280
§7.2 Law Statements Make Explicit Skills	282
§7.3 Implicit Mastery of Techné	295
§7.4 From Techné to Epistemé	300
Epilogue	308
Bibliography	312

Prologue

Philosophy is essentially dialectical. One gets into a dialectic by a puzzle, an *aporia* in thought and understanding. The point of philosophizing is not (necessarily) to get hold of the ultimate, objective, and immutably correct answers. The point is rather to come to be able to see one's way out of the *aporia*, and to understand how one got into it in the first place.

The essay that follows – which I am submitting as my dissertation – is dialectical in two senses. As a piece of philosophizing, the essay is guided by a problem, the problem of understanding how laws of nature are possible, and how it is possible for us to know them. The movement of thought generated by attempts to get out of the problem then yields some ideas that do not stay in the original context in which the problem was felt to exist. Two important ones are, first, perceptual experience is not the only ultimate source of warrant we have for empirical knowledge claims, and second, perceptual experience is not even the only epistemically significant *experience* we can have. These are consequences of the idea that the mastery of skills is a form of interaction with nature that provides epistemic warrant for nomological claims. I shall leave it to the epilogue to examine how this view of skills contrasts with the ways skills are ordinarily thought of in philosophy and the implication of it for empiricism.

The other sense in which the essay is dialectical is more interesting, and it has to do with the way in which I approach the problem that got this essay started, namely, by paying special attention to the dialectic exchange between the realists and the antirealists about the laws of nature. Antirealism about the lawfulness of nature has experienced something like a post-Humean revival since the publication of van Fraassen's *The Scientific Image*. Most, including me, have strong realist intuitions about nomological “connections” in nature. Philosophical positions that are strongly counterintuitive have mostly not ended well in history. So it becomes something of a puzzle why antirealism about laws of nature manages to enjoy popularity every now and then.

Good arguments are needed to establish a strongly counterintuitive position, which, in the case at hand, is a position against naïve realism about nomological connections in nature. What this means is that the antirealists have the burden of argument in the first stage of their dialectic exchange with the naïve realists. If they fail to discharge their

Prologue

burden of argument, or fail to shift the burden of argument back to the naïve realists, we must presume the naïve realists right. This is the sense in which my approach to the problem of laws is highly sensitive of the dialectics of the situation. For I will aim at the arguments with which antirealists try to shift their burden of argument back to the naïve realist. The idea is that, if one succeeds to nip the skeptical dialectic in the bud, then naïve realism is defended simply in virtue of fact that it was the initial default position before antirealists came along. The approach here is *therapeutic*, therefore, in the sense of stopping (philosophical) impulses to follow a particular dialectic, rather than directly combatting (philosophical) behavior.

A successful therapy follows an accurate diagnosis. An accurate diagnosis, in our case, must start with a close look at the arguments with which the antirealists try to move us from the comfort of naïve realism. We soon discover that they have an interesting form. These arguments are *not* directed at laws of nature *per se*. They are *not* arguments to the effect that, for instance, admitting laws into our ontology leads to metaphysical inconsistencies. Instead, the arguments are directed at *our relation as knowers* to putative nomological facts, as well as to *our language use* surrounding law statements.

With the help of a terminology invented by one of the heroes of *inferential semantics*, Dummett, we can characterize the crucial first arguments by antirealists this way. They point out that the way we use nomological statements is *not in harmony*, which means, in the present context, and roughly, that what we treat as evidence for a law statement is too weak for what we treat as admissible consequences of the same law statement. The so-called Problem of Induction, for example, can be formulated in terms of a disharmony between evidence by examination of past instances on the one hand, and the prediction of future instances as admissible consequences, on the other hand.

With this formulation of the antirealists' arguments as the backdrop, the enduring popularity of antirealism begins to make sense. For the arguments are very strong: if there is really disharmony of use, then it seems that the antirealists *must* be right. The familiar example of "tonk"¹ illustrates the point well: no harmony of use, no concept expressed. At the same time, this formulation of antirealists' arguments also shows why various proposals to *analyze* laws of nature in other, presumably safer, more primitive, or better understood terms, all fail to stem out antirealism once and for all. One can *claim* whatever

¹ Prior's invention, „tonk“ is used with the same introduction rule as that of „or“ and with the same elimination rule as that of „and“.

one wants about what “tonk” *really* means, as long as the term’s use pattern is not harmonious, “tonk” will never express a coherent concept. What a naïve realist must now do, therefore, is not to be busy looking for a new analysis of the laws of nature, or refining an existing one. Saying that laws are structural features of possible worlds or relations between universals will not help, if antirealists’ complaint of disharmony of use is not dealt with.

The disharmony charge presupposes a certain picture, a self-understanding, we might say, of the way we go about making law statements. I just mentioned that in Hume’s picture, the evidence side of use essentially consists of observations of instances of co-occurring features, while the consequence side of use includes prediction of future instances of such co-occurrence. New items were added to the consequence side of use: modal and counterfactual statements. But crucially, the evidence side of the picture remained the same. This self-understanding – what I call in the following the *Traditional Picture*² – is shared by both realists and antirealists alike. But since this picture contains disharmony, the realists are in the curious position of fighting battles after having, perhaps unbeknownst to themselves, already surrendered.

This then is the diagnosis of the realism vs. antirealism debate. The *therapy* consists in bringing ourselves to see that the Traditional Picture is wrong, and to see what the *correct* picture is. A law statement is, in the first instance, not backed up by evidence in the form of observation claims. What actually gives warrant for a law statement is the *mastery of a skill through experimentation*. In the simplest kind of cases, the warrant-providing skill has the form: to bring about Y by doing X. Moreover, we can reformulate, though not essentially change, the consequence side of the Traditional Picture in terms of a certain kind of *steps in practical deliberation* that I shall call *desiderative inferences*: I want Y, so I shall X. The harmony of the warrant and the consequence side becomes transparent under this formulation: desiderative inferences are but the preludes to the *exercise* of the skill mastered. As soon as we are able to see our nomological practices this way, we are relieved of the perpetually irritating perception that the concept of law requires, impossibly, moving from finite past observations to claims about the future or contrary-to-fact, non-actual, events.

² C.f. §4.

Prologue

The therapy that results, if successful, has implications beyond the realism debate about laws of nature. One radical rethinking required is the recognition of a second basic mode of epistemic access we have to the world around us next to perception, in the form of *skills*. That then leads to a revised concept of *experience* that figures in an intelligent empiricism. A proper treatment of a substantially modified empiricism in the Kantian tradition, however, must await another occasion.

The individual parts of the essays are organized as follows. **Part I: Semantic Preludes** contains two chapters purely on semantic theories. §1 is devoted to possible worlds, which are widely thought to be relevant to modal contents. There I distinguish semantic explanation from semantic modeling, and show how possible worlds come up as useful heuristics in semantic modeling tasks. In §2 I review a semantic tradition that explicitly aims to be explanatory, namely, semantic functionalism. My review is selective because I almost exclusively discuss *inferential* semantic functionalism, and focus on three authors: Sellars, Dummett, and Brandom. I distinguish more modest forms of semantic explanation from more demanding inferentialist programs that aim to explain linguistic *understanding* and meaningfulness as such, and that are thereby inextricably tied up with complex questions about norms and the philosophy of mind. Much of the chapter is aimed at clarifying the kind of semantic question I would be interested in in connection with nomological statements.

In **Part II: Semantic Functionalism & The Problems of Laws** I explore ideas from semantic functionalists that are relevant to laws of nature. I find, in §3, that none of the three major semantic functionalists is able to respect the naïve factualism about laws of nature statements. But I also note a promising idea, found in Brandom, that some claims *make explicit* what I call proto-contents implicit in a doing or experience. Although this idea does not directly help defeat nomological anti-realism, it points at a suggestive epistemological picture for laws. In §4 I explore the idea of Harmony, and apply the idea to diagnose the permanent impasse between realists and antirealists about laws.

Part III: Laws and Skills – A Therapy describes the *correct* picture of language use surrounding laws of nature statements. Throughout this part I assume a simple form of law statement: what I call *binary* law statements. §5 describes the basic elements of a nomological practice: skills acquired by experiments, and steps in practical deliberations I call *desiderative inferences*. §6 deals with a number of potential objections, including those

about perception and status of empiricism, those about the nature of non-perceptual elements in the picture developed in §5, and skepticism.

In **Part IV: Beyond Therapy** I apply the idea of explicit-making to give a kind of semantic explanation for law statements, which amounts to completing the account of the epistemology for laws. Appeal to the explicit-making of nomological information is now intelligible, *given* that **Part III** is successful in stopping the skeptic dialectic and in reinstating the presumption of naïve realism about laws. I then extend the picture of nomological practice developed in §5, to cover those practices that lead to complex scientific law statements. Also in this second part, the idea of *explicit making* will play a central role.

Part I: Semantic Preludes

§1 Possible Worlds and Semantic Modeling

§1.1 Semantic Modeling vs. Semantic Explanation	10
§1.2 Possible Worlds I – Intensional Semantic Modeling.....	11
§1.2.1 Extensional Semantic Theory	12
§1.2.2 Compositional Modeling.....	13
§1.2.2.1 Compositionality of Meaning vs. Compositional Modeling.....	13
§1.2.2.2 The Empirical Discovery View vs. The Modeling View	14
§1.2.2.2a Meaning Without Assignments	14
§1.2.2.2b Within The Assignment Paradigm.....	15
§1.2.2.2c Disputes About Assignments.....	16
§1.2.2.2d Janssen’s Argument For Compositionality.....	18
§1.2.3 Fine-grained-ness and Intensional Semantic Modeling.....	19
§1.3 Conceptions of Possible Worlds	22
§1.3.1 Lewis: Parallel Materially Existing Worlds.....	22
§1.3.2 Kripke: Stipulative Understanding Of Possible Worlds	24
Appendix to §1.3.2: Cognitive Significance as Meaning	26
§1.3.3 Possible Worlds As Representations.....	29
§1.3.3.1 Possible Representations, vs. Representations of Possibles	29
§1.3.4 Possible Worlds as Heuristics For Semantic Modeling	32
§1.3.4.1 Epistemic vs. Explanatory Priority	33
§1.3.5 Which Sense of “Possible”?	35
§1.4 Possible Worlds II – Modal Vocabulary.....	36
§1.4.1 Non-Extentional Contexts Generating Operators	36
§1.4.1 Variety and Iteration	38
§1.4.1a Variety of Operators.....	38
§1.4.1b Iteration of Operators.....	39
§1.4.3 Non-Modal Conceptual Possible Worlds Are Most Basic.....	40
§1.5 Conclusions.....	43

§1.1 Semantic Modeling vs. Semantic Explanation

This essay deals in, among other things, semantics, to the extent that it is about certain types of modal semantic content: that of a law of nature statement. But it does not deal equally with different *aspects* of semantics. The first distinction to be drawn, then, is between two kinds of semantic projects, one *explanatory* of meaning, the other model-theoretic, in the sense of meaning *modeling*.

The distinction between semantic modeling and semantic explanation is not an idiosyncrasy of mine. Many philosophers who work on model-theoretic semantics freely acknowledge that their theorizing is not meant to explain what linguistic meaning consists in, that it is to be distinguished from “metaphysical semantics”, where the *ti esti*, or nature question about linguistic meaning is presumably answered. Some even go so far as to say that what constitutes an expression’s having a determinate meaning is not a semantic topic, but that of a general “philosophy of language”. Here is how Ned Block puts it in his entry on conceptual role semantics in *The Routledge Encyclopedia of Philosophy*:

There are two quite different projects that go by the name 'semantics'. One, which we might call *linguistic* semantics, deals with the meanings of particular expressions in particular languages and how they fit together to make up meanings of larger expressions, for example, how the meaning of 'John hopes angels exist' is related to the meaning of 'angels exist'. The second project, *metaphysical* semantics, is one of investigating the fundamental nature of meaning, especially what it is about a person that gives his words or thoughts whatever meanings they have in the first place.

Metaphysical semantic theories attempt to deal with questions such as semantic holism (q.v.) and whether meaning can be specified independently of a thinker's environment (see Content: Wide and Narrow). Examples of such theories are *causal theories of reference*, *teleological* (q.v.) theories that try to explain meaning in terms of evolution, and *informational* (q.v.) theories that construe meaning as some sort of covariation relation between contentful states of the person and the phenomena in the world that those states pick out. Linguistic and metaphysical semantic theories are largely independent. For

example, the difference between teleological and informational theories has no impact on any issue in linguistic semantics.

Given this distinction of modeling vs. explanation, the primary aim in this essay is *not* semantic modeling, but a kind of semantic explanation. Later in the essay, in §2, further distinctions will be made between various types of semantic explanation. For now note that my focus on a kind of *modality*, as opposed to say truth and reference, gets me perforce in close quarters with the infamous *possible worlds*.

Now possible worlds have been invoked in *both* the explanatory and the modeling project, though frequently it is not explicitly stated for *which* of these goals they are called upon. My view of possible worlds is that they are to be thought of as heuristics in semantic modeling. Given the depth of the influence of such modal realists as David Lewis, who treated possible worlds as both explanatory and good for modeling, it is appropriate to devote some space to describe with some precision the limited role I believe possible worlds can play in semantics.

In this chapter, then, I develop the idea of possible worlds from the ground up, in the context of the project to develop a systematic representation, or modeling, of linguistic meaning. The thought is that once we develop the ideas carefully from their roots, a lot of false impressions of and expectations from possible worlds will fall of themselves. The thesis I want to maintain is that semantic modeling, when done with sufficient sophistication or for relatively sophisticated languages, *requires* something that we might as well call “possible worlds”. But the phrase “possible world” plays a merely suggestive, heuristic role in guiding the construction of abstract modeling schemes, rather than designating entities of a special kind. To develop the ideas in their clearest forms, I shall proceed incrementally, beginning with semantic modeling for simple, non-modal languages.

§1.2 Possible Worlds I – Intensional Semantic Modeling

Consider a typical view on how the *non-explanatory* sort of semantic theory aims to do in general (underlines mine):

One task is simply that of explaining how signs and symbols can ever be about anything at all, how they come to *have* semantic content. This is not generally considered part of the field of semantics proper, but it is the topic of closely related

Part I: Semantic Preludes

work in the philosophy of language. Rather semanticists presuppose that symbols somehow acquire content and take their task to be that of characterizing the semantic content of arbitrary complex expressions, given the content of their basic constituents. This is usually broken into two related parts. One is to describe the meanings of the semantically basic items of the language. This is called *lexical semantics* and deals with the meanings of expressions such as “friend” and “my.” The other part is to give a precise characterization of how the meanings of complex expressions depend on the meanings of their constituent parts. This is called *compositional semantics* and deals with the meanings of expressions such as “my friend” and “my friend’s friend won.”

– *Jon Barwise and John Etchemendy “Model-Theoretic Semantics”*

The idea is that model-theoretic semantics tells us how a semantic fact about a complex expression, say a sentence, depends on semantic facts about re-combinable elements (e.g. words) that make up this complex expression and the way these elements are combined. Let me illustrate this idea with examples. In the course of doing so, a picture of semantic modeling will slowly take shape.

§1.2.1 Extensional Semantic Theory

A simple (extensional) model-theoretic approach to the sentence “Fido yawns” would involve assigning Fido to “Fido”, the set of yawning entities to “yawns” and a truth-value to the entire sentence depending on whether it is true or false. In general, an object will be assigned to a singular term, a function from objects to truth-values to a predicate expression and a truth-value to a sentence. These assignments taken collectively *model* the way the meaning of a sentence depends on the meaning of its simple syntactic components: the sentential assignment is **True** just when the assignment to the predicate expression yield **True**, when applied to the assignment to the subject term. Compositionality of meaning is modeled, in other words, by the formal operation of functional application. To give a slightly more complex example, consider the determiner “all”. The semantic approach described assigns a function from pairs of sets to truth-values, formally: $\lambda A \lambda B (A \subseteq B)$. The compositionality of meaning for a sentence like “all dogs are friendly” is then modeled by two successive functional application of [all], first on [dog], then on [friendly], where the square brackets mean “the assignment to”.

§1.2.2 Compositional Modeling

§1.2.2.1 Compositionality of Meaning vs. Compositional Modeling

Before going into considerations that prompt refinements to this simple semantic theory, let me examine the sense in which the mini-theory I just sketched *models* meaning. The first distinction to be made is between the *aspect of our language to be modeled* from *the way we model that aspect*. One target aspect, as we saw from Barwise and Etchemendy, is “how the meanings of complex expressions depend on the meanings of their constituent parts”. *That* there is such a dependence (not that the dependence is complete, or that it has this or that particular form) I take to be established by the learnability and productivity considerations. But it is important to distinguish this *highly unspecific and amorphous* claim about our language – call it the **compositionality of meaning** – from a principle of modeling: the requirement that the assignment scheme that lies at the center of a modeling project satisfy certain *formal properties*, such as “the semantic assignment to the concatenation of a predicate and subject is identical with the *result of functional application* of the assignment to the predicate to the assignment to the subject”. Let us call principles of the latter kind **Principles of Compositional Modeling**. The principle of compositional modeling for the mini-theory sketched, can be formulated more formally::

(One) **Principle of Compositional Modeling**: If ‘*a*’ is a singular term, ‘*P*’ is a predicate, ‘*[]*’ the semantic assignment function, then the following holds:

$$[P(a)] = [P]([a])$$

This principle of compositional modeling is *one way* to model the non-specific factum of compositionality of meaning. Principles of compositionality modeling are *precise*, they are *chosen*, and they depend on the overall modeling projects for which they are chosen (e.g. an intensional modeling scheme like that of Montague³ has a slightly different principle of compositional modeling). By contrast, the compositionality of meaning is vague, confirmed by pre-theoretic intuition and such general arguments as the learnability of language.

³ C.f. Montague’s seminal essay from 1970 “Universal Grammar”. Another locus classicus is his “English as a Formal Language” from the same year.

Part I: Semantic Preludes

Compositionality of meaning is not the only aspect modeled. Other aspect can include: (pre-theoretic) notions of reference, truth, and entailment relations. I shall illustrate semantic modeling primarily from the perspective of their attempt to capture compositionality of meaning.

§1.2.2.2 The Empirical Discovery View vs. The Modeling View

With the distinction just made, I have adopted, without ceremony, what I shall call the “The Modeling View”: the view that *precise* principles of compositionality is a *regulative principle* for semantic modeling, rather than *semantic facts we discover*. The contrasting view I will call “The Empirical Discovery View”. It is the idea that *after* we, with help of a semantic theory, determine the meaning of singular terms to be objects, that of predicate expressions to be sets of objects, that of sentences to be truth-values, *we then discover* that the meaning of a sentence depends systematically on the meanings of its subject and predicate terms in some precisely describable manner: in fact, describable by the functional application.

§1.2.2.2a Meaning Without Assignments

The basic reason why the empirical discovery view is implausible is that there is no universal agreement at all as to what kind of thing the meaning of an expression is, or even whether a statement describing the meaning of an expression should have the form of relating that expression to another entity, that is, whether it should have the form of an *assignment scheme*. There are two dimensions of variation that are relevant here. On the one hand, there are *external* alternatives, that is, alternatives to the assignment-cum-functional application paradigm. One such alternative paradigm is a version of Davidsonian truth-conditional semantic theory propounded by John McDowell and Gareth Evans, according to which the meaning of an expression is *displayed* by a meaning-description, rather than referred to and then correlated with that expression (i.e. ‘The meaning/semantic value of X is N’, where ‘N’ is a nominal expression referring to the meaning of X). The family of semantic theories under the rubric “use-theoretic semantics” represents another alternative. In these theories meaning-descriptions neither refer to nor display meanings. They are rather specifications of how expressions are, or ought to be, used. So the first dimension of variation concerns the very *form of meaning-descriptions*. Assignment-cum-functional application paradigm is only one amongst various possibilities.

§1.2.2.2b Within The Assignment Paradigm

The second dimension of variation is *within* the family of theories exemplifying the assignment-cum-functional-application paradigm. There is no consensus, nor is a consensus generally thought to be necessary, about what kind of assignment for one category of expressions goes with what kind of assignment for another category of expressions. Frege himself devised two kinds of assignments, for all categories of expressions, one he calls the *Bedeutung* of an expression, the other he calls its *Sinn*. A Fregean holding the Empirical Discovery View would think – though it is hard to determine what Frege himself thought – that *Bedeutung* and *Sinn* are *natural kinds* that each exhibits, independent of us semanticists, a compositional structure which we then discover. Such a person would think that *truth-values* to be the kind of sentential assignment that *naturally* goes with the assignment that relates a singular term to the object it refers to, and both go, again, naturally, with the assignment to a predicate expression its extension – the set of objects falling under it.

The dominating influence of, and the incredible stability enjoyed by Frege's semantic theory until this day creates the powerful illusion, but an illusion none-the-less, that there are such semantic natural kinds as sense and reference across all categories of expressions. Call this the Doctrine of Semantic Natural Kinds. The best medicine against this doctrine is a survey of the internal variation I have been speaking about.

Before turning to that, consider for a moment what Frege actually had to say on this topic. His considerations for assigning truth-values as “*Bedeutung*” to sentences, are, when considered as arguments for a semantic natural kind across all syntactic categories, not very persuasive. But these considerations were also not intended by Frege to be such arguments. So, when introducing the idea in his seminal article *Über Sinn und Bedeutung*, Frege advances the following consideration in support of his alignment of truth-values for sentences with reference for singular terms. First of all, a sentence's having sense (“*Sinn*”) can be ascertained independent of the question whether a name occurring in that sentence really refers or not. The lack of reference of a singular term contained in a sentence, therefore, need not mean the lack of sense for the sentence. On the assumption that a sentence *has* a reference, and the reference of the sentence depends on the reference of its simpler constituents (this is the “principle of compositional modeling” for the particular assignment

Part I: Semantic Preludes

scheme with “Bedeutung”), Frege concludes that whatever the reference of a sentence is, it cannot be its sense.⁴ On the other hand, Frege notes, the question whether a name refers becomes relevant if, and *only if*, we are interested in finding out if the sentence containing that name is true or false.⁵ This line of thought purports to show that *given* that “Bedeutungen” are to be assigned to sentences, truth-values ought to be assigned. This is a conditional claim: its recommendation is conditioned on *our desire* to have a global assignment scheme meeting certain principle of compositional modeling. Whatever one thinks of the merit of this argument by Frege, it is not even *intended* as establishing the *un*-conditional claim that there is a natural semantic category, “reference”, both for sentences and for sub-sentential expressions. A similar claim has been made again, later, by Church, on the basis of one version of the so-called “slingshot arguments”⁶. But again, there has been no serious argument for the *un*-conditional claim that sentences have reference, above and beyond the mere practical necessity that a theory of meaning forged in the assignment-cum-functional application style *has* to assign *something* to sentences the way it assigns references to singular terms. If anything, the fact that some versions of the “slingshot argument” forces us to make the perfectly *unnatural* conclusion that sentences *refer to truth-values*, in a sense of “refer” that is supposed to be nontechnical, speaks against the assumption that the *technical* notion of reference is a natural category across expression types.

§1.2.2.2c Disputes About Assignments

As the next step towards disposing of the Doctrine of Semantic Natural Kinds, consider the internal disagreements within the assignment-cum-functional-application paradigm. While Frege thought that the assignment of the referent to a singular term

⁴ It is also instructive to note that, in these paragraphs, there is an ascent, as it were, from an intuitive syncategorematic understanding of the phrase ‘has sense’ as simply meaning ‘meaningful’, to the technical employment of the term of art ‘sense’ as denoting something that can be had by an expression and specific to the expression having it. This helps Frege to make plausible his sense-assignment-scheme.

⁵ ‘Der Gedanke bleibt derselbe, ob der Name "Odysseus" eine Bedeutung hat oder nicht. Daß wir uns überhaupt um die Bedeutung eines Satzteils bemühen, ist ein Zeichen dafür, daß wir auch für den Satz selbst eine Bedeutung im allgemeinen anerkennen und fordern.... Warum wollen wir denn aber, daß jeder Eigenname nicht nur einen Sinn, sondern auch eine Bedeutung habe? Warum genügt uns der Gedanke nicht? Weil und soweit es uns auf seinen Wahrheitswert ankommt.’ – *Über Sinn und Bedeutung*, pg. 33.

⁶ C.f. Church, Introduction to Mathematical Logic, pg. 24-5.

is aligned with assignment of a truth-value to a sentence, this has been disputed in recent times. Alternative assignment systems have been developed in which reference-assignment (for singular terms) corresponds to *state-of-affairs* assignment to sentences, rather than truth-value assignment (c.f. for example, Graeme Forbes, *Languages of Modality*, pg. 131 ff.).⁷ What gives Forbes the conceptual opening for his alternative choice of sentential references is what Kripke's work reveals to be a profound difference between two kinds of sentential contexts treated equally by Frege: modal contexts on the one hand, and propositional attitude report contexts on the other. Frege thought that what distinguishes *both* sorts of contexts is that, for sentences embedded in these contexts, the ordinary (free-standing) sentential *Sinn* plays the role of component sentential *Bedeutung*. The difference between these two kinds of context might be known before the work of Kripke, but it only became decisive after Kripke showed convincingly that *cognitive significance*, the cornerstone of the Fregean notion *Sinn*, need not contribute to the *modal behavior* of an expression. The modal contexts are sensitive only to the modal behavior of the sentences imbedded in them, not to their cognitive significance.

A Fregean can react to Kripke's insight either of two ways. She can say that the concept of cognitive significance is irrelevant to *meaning*, and that only the modal behavior of an expression is. Or she can say, with Frege, that cognitive significance is the ultimate measure of meaning, but the concept of *reference* needs to be richer and to reflect the modal-behavior of an expression (when the expression is a proper name, we need not change anything, because the reference of a proper name is rigid modally speaking). The latter kind of reaction is precisely what Forbes' alternative proposal amounts to. Forbes' "states of affairs" are designed so that their assignment to sentences captures *just* the modal behavior of these sentences, but not their cognitive significance. Frege's strategy of invoking the *Sinn* of imbedded sentences for treating sentential contexts that sensitive to differences not registered by free-standing reference, is not affected. But the only contexts for which Frobes would *need* to apply the strategy are attitudinal contexts.

⁷ Forbes' reasoning behind his choice of reference for sentences reveals other, complex problems with the Fregean treatment of attitude and modal contexts. I shall go into these here.

Part I: Semantic Preludes

Collectively, these considerations⁸, the dispute about the nature of meaning, the lack of a substantive argument in Frege's original writings, the variations and dispute within the assignment paradigm, etc., put in doubt the Doctrine of Semantic Natural Kind, the thesis that there is a natural semantic kind available for *all* classes of expressions. If such a semantic kind existed – one to which, let us say, both reference and truth-value belonged – one *would* be able to say that the assignments *are* just the formal articulation of these semantic facts that semanticists *discover*. What appears to be the case, instead, is that the assignments are subject to the *regulative principle* that some formal criterion of compositionality should be met. What form such a criterion for compositional modeling takes, depends in large part on which assignment scheme *we* want to adopt.

§1.2.2.2d Janssen's Argument For Compositionality

Arguments about compositionality fall into two categories. One kind of arguments aim to establish or refute the very *desirability* of having a theory of meaning that is subject to some principle of compositional modeling. An example of such an argument would be an argument from the learnability of any natural language. These arguments do not establish the precise form of the principle of compositional modeling, only the desirability of having one. Another kind of arguments are about the technical feasibility of representing certain linguistic phenomena compositionally within a certain theoretical setup, for example the assignment-cum-functional-application framework. Examples of the latter kind of arguments are Higginbotham's suggestion⁹ that “unless” does not behave compositionally when it occurs in certain combination with quantifiers, as well as Pelletier's reply to that¹⁰. Theo Janssen recently argued convincingly that there are not enough *technical limits* to the complexity of a model-theoretic representation of meaning to force us to abandon the principle of compositional modeling.¹¹ Janssen does that by showing that various

⁸ Note also, assuming that truth-value assignment for sentences go naturally with referent-assignment for singular terms, the question about what the assignment for predicate expression remains open until we enforce, in addition to compositionality principle, some kind of context principle, to the effect that the assignment should not more finely distinguish predicate expressions than is necessary to account for the differences in assignments to sentences in which these predicate expressions occur.

⁹ Higginbotham 1986, Linguistic theory and Davidson's program in semantics.

¹⁰ Pelletier 1993, On An Argument Against Semantic Compositionality.

¹¹ C.f. Janssen 1997, *Compositionality*.

alleged counterexamples to compositionality of meaning – that is, examples where aspects of language are thought to be incapable of being modeled in accordance with some principle of compositional modeling – can be easily removed by incorporating new variables into a model-theoretic semantic theory. The question, in other words, is not whether the language we want to model *permits* an assignment scheme conforming to some principle of compositional modeling. The only question is whether *we want* to so model a language. The answer to *that* question, depends on whether we find what I have called the “intuition” about the compositionality of meaning to be plausibly established by a learnability argument.

Considerations of *technical feasibility*, such as those advanced by Janssen, complement the argument for the *desirability* of devising semantic theories in accordance with some formal principle of compositional modeling. It is good to know that what is good to have, *can* also be had. Henceforth I shall assume the semantic modeling under discussion to come with some formal principle of compositionality.

§1.2.3 Fine-grained-ness and Intensional Semantic Modeling

Time to return to our story about possible worlds in model-theoretic semantics. One clear manifestation of the fact that the simple semantic theory I sketched above is merely a *model* for the compositionality of meaning is that the assignments are not nearly fine-grained enough to capture what we intuitively grasp as differences in meaning. So “Fido” and “Donald’s dog” may have the same referent – so the same assignment under the extensional scheme – but, as is familiar from some famous arguments by Kripke, they cannot have the same meaning. So, had Donald decided to buy his dog from a professional breeder as opposed to picking one up at the local animal shelter, Fido would *not* have been Donald’s dog. Consequently, the identity of reference between “Fido” and “Donald’s dog” is the result of an accident in the world, rather than the sameness of meaning of these two phrases.¹² Moreover, it is conceivable that we discover that Fido is not Donald’s dog after all, but is merely

¹² I am appealing to Kripke’s argument that names do not have descriptive meanings here. I have not made the point by reference to Frege’s distinction between sense and reference, in order to avoid the dispute over whether the meaning of a name is just its reference. For more on this debate, see Scott Soames’ *Beyond Rigidity*.

An easier way of making the same point without invoking Kripke and the intricacies of proper names, as Ede Zimmermann pointed out, is to consider two definite descriptions that happen to refer to the same thing, e.g. “Donald’s dog” and “The gift Donald received from his sister last Christmas”.

Part I: Semantic Preludes

being looked after while its master is on vacation. That discovery would not be a case of change in meaning for the phrases “Fido” and “Donald’s dog”.¹³ Assigning the referent to a singular term therefore fails to register differences in meaning.¹⁴ Assigning truth-values to sentences similarly fails to reflect meaning differences among all the true sentences as well as differences among all false sentences. Finally, it is a familiar story told over and over again in semantic textbooks that the extensions of predicate expressions can coincide even if they are wildly different in meaning. So it may happen that everyone who is sleeping is also snorting, but that does not mean that “is snorting” and “is sleeping” are synonymous predicates. To the extent that extensional assignments *abstract* from finer differences of meaning over and above reference and truth, therefore, they provide an over-simplified model for meaning and its compositionality.

The defect of the extensional assignment under discussion exists not only on the level of “semantic intuitions”. It takes a very concrete form of *failure to model an aspect of language use*: the assignment scheme fails to model *entailment relations*. The identity of the extension of “is sleeping” with that of “is snorting” does not tell us which of the two sentences entails the other: “John is sleeping” or “John is snorting”.

The issue has been called the fine-grained-ness problem for extensional semantic theories like the one I have sketched here¹⁵. The problem can be roughly diagnosed as follows. The truth of a sentence like “Fido is bored” and the reference of a singular term like “Donald’s dog” each depends on two factors: the meaning of the

¹³ These two arguments, modal and epistemological respectively, were advanced by Kripke against two versions of a descriptive theory of meaning for proper names. The modal argument maintains that proper names, unlike definite descriptions, refer to the same object in all counterfactual situations. The epistemological argument rejects a weakened descriptivist thesis: a definite description determines the actual *referent* of a proper name. The conclusion of the first, modal argument, namely proper names are rigid designators, found almost universal acceptance, whereas the conclusions of the second, epistemological argument have been frequently challenged. So it has been proposed that the referent of a proper name is fixed by a description after all, for – some alleges – Kripke’s own account of reference fixing in terms of baptism and the causal chain of use is nothing but such a reference-fixing description. More generally, the Kripkean idea that the meaning of a proper name is exhausted by its modal behavior has been controversial, though not without its defenders (e.g. Scott Soames, *Beyond Rigidity*). Two-dimensionalism, for example, has been a popular anti-Kripkean paradigm, in which semanticists attempt to resurrect *some* version of cognitive significance or reference-fixing feature as part of the *meaning* of a proper name. See more discussion below.

¹⁴ One might as well say, with Frege, that reference does not determine sense. But I cannot put the matter that way because I am try to develop something like a sense-assignment scheme *out of* the extensional scheme.

¹⁵ C.f., for example, *Model-Theoretic Semantics* by Jon Barwise and John Etchemendy.

words making up the sentence or the singular term on the one hand, and the way the world is. A systematic assignment scheme that tracks only truth and reference, therefore, is not guaranteed to register all differences in meaning, as differences in meaning can be erased after the contribution of the world-factor is counted in. The three kinds of examples just mentioned (e.g. proper name and definite description, all true sentences, and two co-extensional predicate expressions) show that this actually happens, systematically and on a massive scale. Tracking truth and reference alone, therefore, is too “coarse-grained” a requirement for modeling meaning, because truth and reference reflect the *combined* effect of meaning of words and the way the world is.

This is the place where more complicated assignment systems, specifically the so-called *intensional* semantic theories, going back to Carnap, Church and possibly the young Wittgenstein, are motivated and introduced. Therein also lies one origin of what one now calls *possible world* semantics. The animating thought for this large family of semantic theories can be put this way: to correctly model meaning, an assignment scheme must not only track *actual* truth and reference, but also how truth and reference *depend on the ways the world might be*. If an assignment scheme can model *that*, then it can obviously model meaning more accurately than an extensional scheme, on account of the examples just discussed. Intensional semantics, to put it very concisely, can be thought of as just such a project: to work out a ***systematic assignment scheme that models the dependence of linguistic expressions for their truth or reference on the ways the world might be.***

To get a taste of what an intensional assignment scheme might look like, suppose we represent all the different ways the world might be as a set W (leaving aside for the moment the question about their ontological status). An assignment scheme designed to track truth and reference given any way the world might be would assign to expressions various *functions* with the set W as domain. To continue with our example above, the assignment to “Fido” would be a function from W to objects, the assignment to “yawns” a function from elements of W to sets of yawning objects in the respective worlds (or functions mapping yawning objects in the world to **True** and objects that are not yawning to **False**), and finally the assignment to the sentence “Fido yawns” would be a function mapping possible worlds relative to which the sentence is true to **True** and the rest to **False**. The combination of assignments is now

more complicated: [Fido][yawns] (w) = [yawns] (w) ([Fido](w)). This illustrates a somewhat different principle of compositional modeling than the one we had for our extensional theory. Still, functional application is at the center of it.

Let me call elements of *W*, namely, the different ways the world might be, “possible worlds”. We shall use the term as a placeholder until we eventually find a more concrete description to fill it. Note that the way we have come to the idea of possible worlds does not require the presence of *modal expressions* in the language for which we want to do semantic modeling. Before discussing those, I will first turn to the question: what are, after all, possible worlds?

§1.3 Conceptions of Possible Worlds

§1.3.1 Lewis: Parallel Materially Existing Worlds

The first conception under review, then, has it that possible worlds exist in the same way the actual world – the only one a philosophically naïve point of view would recognize – exists, presumably causally isolated from each other and from the actual world. This in my view is not a plausible position. But since the task of this chapter, as far as model-theoretic semantics is concerned, is to show the possibility of a philosophically uncontroversial role for it, rather than to argue that that role is the only sensible one it can take, I will be rather brief and mention only the main problems of the modal-realist position.

First, there is simply no evidence for parallel existing worlds. There is a pattern of arguments for postulating questionable entities having the form: a semantic theory assuming entities of type *X* (e.g. possible-worlds in the sense of worlds existing parallel to the actual one, or merely possible objects) can deal with languages with greater expressive power than semantic theories that do not assume the existence of these entities; *ergo*, it is better to assume entities of type *X*.¹⁶ These arguments suffer from a fundamental confusion. A fact of the form “*a practical goal can be better achieved* if a certain assumption is true than if the assumption is false” has no *evidential* bearing on that assumption. If it were true that nuclear fission generates no radioactive waste, it would be much easier to solve the world’s energy problem. This conditional fact of expediency, however, is no *evidence* for the truth of the

¹⁶ C.f. Forbes’ *Languages of Modalities* for an example for merely possible objects. Lewis’ *Plurality of Worlds* is also full of such arguments.

assumption about nuclear fission. Expediency is not the native soil for evidential relations.

Second, methodologically, there is absolutely no motivation for postulating worlds that exist in the same sense this world exists. Remember that we are aiming for a *model* for meaning built around an intensional assignment scheme. For *that* purpose, it is absolutely immaterial what we use to do the modeling. Certainly, abstract entities can do just as well as things having material existence. But let us *suppose* for the sake of the argument that one wanted to *explain* what meaning consists in by reference to possible worlds, presumably because one thinks that “sleep” means what it means *because* the word classifies the objects it classifies in various possible worlds (not just we can *tell* its meaning by thinking about what it would classify had the world been different).¹⁷ Even in that case, it is not clear that one needs a modal-realist conception of possible world. It might be enough to postulate possible worlds as *non-material* entities, standing in a similar relation to the actual world as various properties stand to an object that might bear them. One might put this last point by saying that the view that possible worlds, rather than modal facts, have explanatory priority needs not be committed to the *material* existence of worlds parallel to ours.

Third, it will be very difficult to give an account of how we come to *know* facts about possible worlds as conceived by modal realists. Remember that we introduced possible worlds because they enable us to *know and demonstrate* very quickly that two co-extensional predicates differ in meaning.¹⁸ But if possible worlds had causally isolated material existence, it would be very puzzling how we could rely on possible worlds to facilitate our knowledge about and awareness of meaning, as we in fact seem to do. Critiques of modal realism on account of its bizarre epistemological consequences include for example Van Fraassen’s complaints against necessitarianist account of laws, namely, if one tries to explain laws in terms of structural features of possible worlds (understood as the modal realists understand them), then it becomes unclear how these laws could possibly have implications about *this world*¹⁹.

¹⁷ See discussion below in §1.3.4.1 Epistemic vs. Explanatory Priority.

¹⁸ See §1.3.4.1 Epistemic vs. Explanatory Priority below for more details.

¹⁹ Cf. *Laws and Symmetry*, especially the section on branching-time model for objective chance.

Part I: Semantic Preludes

For all these three reasons – lack of evidence, lack motivation, and problems for modal epistemology – and because of the methodological orientation of adopted here, I will leave modal realism behind for the rest of this essay.

§1.3.2 Kripke: Stipulative Understanding Of Possible Worlds

In his milestone lectures *Naming and Necessity*, which was read out more than three and half decades ago, Kripke formulated and defended his signature claim that proper names are “rigid designators”. These lectures, and this claim in particular, have had profound implications in many areas of philosophy. In semantics, two exemplary influences of Kripke’s lectures are: (i) putting in doubt the tight connection, often traced to Frege, between *cognitive significance* and *meaning*; (ii) setting forth a striking understanding of the concept of possible worlds - probably *the* most fundamental concept to modern semantics.²⁰ The first of these points I shall discuss in the Appendix that follows. As for the second point, Kripke’s essays laid out a conception of possible worlds that is diametrically opposed to Lewis’ modal realism just rejected. The rest of the section shall be a review Kripke’s *positive* view about possible worlds.

The relevance of possible worlds in Kripke’s lectures is manifested in his diagnosis of the *source* of rigidity of proper names. That diagnosis is based on what we might call a *stipulative understanding* of possible worlds. In *Naming and Necessity*, Kripke had the following to say about possible-worlds:

‘possible worlds’ are stipulated, not discovered by powerful telescopes
– p44. Naming and Necessity (henceforth NN)

Let me call this suggestive but as yet vague formulation the Principle of Stipulation. Kripke elaborates how this has to do with rigidity of reference in the same lecture:

Those who have argued that to make sense of the notion of rigid designator, we must antecedently make sense of ‘criteria of transworld identity’ have

²⁰ *Naming and Necessity* also stand in the beginning of a tradition that was to unfold into a movement of sorts identified by most now as “semantic externalism”, a movement that is frequently associated with Putnam’s seminal Essay *The Meaning of Meaning*, which appeared three short years after Kripke’s lectures. Kripke’s semantic externalism consists in rejecting – in the so-called “epistemological arguments” (see the main text below) – the view that the reference of a proper name is determined by some *a priori* knowable description, and replacing it with a so-called “causal theory” of how the reference of a name is determined.

precisely reversed the cart and the horse: it is because we can refer (rigidly) to Nixon, and stipulate that we are speaking of what might have happened to him (under certain circumstances), that 'transworld identifications' are unproblematic in such cases. – p49. NN

we do not begin with worlds... and then ask for criteria of identification of objects across worlds; on the contrary, we begin with objects, which we have, and can identify, in the actual world. We can then ask whether certain things might have been true of the objects. – p53. NN

So if we accept Kripke's analysis of why proper names are rigid designators, we would seem to be committed to a particular way of understanding possible worlds and functions there-from. And even if it is not quite clear yet what this understanding looks like in detail, it clearly represents a rejection of the Lewisian realism about possible worlds.

Just what the "stipulative" understanding of possible worlds amounts to is not so straightforward to say as it can seem at first sight. Surely I cannot stipulate a possible world by decreeing any combination of states of affairs whatsoever to be possible, much like certain processes in Congress, by virtue of being the kind of processes they are, yield laws. The problem is not that I may not have the right authority to make the decree. *No one* can decide that it is a possibility that Columbus was both the first and the third European to have sailed to America, or that there is a triangular square. These are simply not possibilities. So whether something is stipulatable – in the sense that the stipulation will be stipulation of a *possibility* – is a question to which, it seems, there can be an objective answer. Moreover, if possible worlds are simply creatures of our unconstrained will to stipulate, then they cannot serve as the basis of a model for meaning. Remember that the initial idea of an intensional assignment scheme was that, one way to tell that "sleep" and "snort" are distinct in meaning is to observe that the world might be such that one but not the other is truly said of John. If possibilities could be freely stipulated, then for *any* two predicates of different morphological form, say A and B, one can "tell" that they are distinct in meaning by observing that in *one* possibility at least – namely the one we *stipulate* to be such that John is A but not B – one predicate but not the other is truly said of John. So quite independent of the question whether possibilities in fact are objectively determined, *if* they could be freely stipulated, they would not serve the

Part I: Semantic Preludes

purpose of modeling meaning, which is the only reason why they were invoked in our intensional assignment scheme in the first place. In short then, the stipulative understanding of possible worlds, as a component of our semantic model, cannot be a view having the consequence that, when it comes to possibilities, saying so makes it so.²¹

One way to plausibly construe Kripke's stipulative understanding of possible worlds, is to replace "stipulations" with "representations" of possibilities. This is the next conception I now turn to.

*Appendix to §1.3.2: Cognitive Significance as Meaning*²²

The Fregean theory of *sense* – which has become a mighty tradition in recent semantic work – is a concept closely tied to the term of art "cognitive significance". Against that tradition, Kripke's argument that proper names are rigid designators shows that it is possible for cognitive significance to outstrip what can be represented by the dependence of reference and truth on ways the world might be (or, possible worlds). So if we take "possible ways the world might be" to mean what Kripke means by "metaphysical possibilities", then difference in cognitive significance between two co-referential proper names will not be detectable in the ways the names' reference depends on the ways the world might be. This means that a separate decision must be made, at least for these expressions, about whether we want to model cognitive significances going beyond dependence for truth and reference on worlds. If we, following the *basic* intensional semantic program, decide that that dependence is an adequate representation of meaning, then we would be adopting a notion of meaning according to which cognitive significance need not be part of it.²³ More will be said about the issue shortly. But whichever side wins the debate on the issue of *extra* fine-grainedness in terms of cognitive significance, it does not affect the program of constructing a systematic intensional assignment scheme as a second-order approximation, just as the fact that *actual* truth and reference are not fine-grained enough to adequately represent meaning does not detract the status of extensional assignment schemes as a *first-order* approximation.

²¹ The semantic argument in this paragraph – that possible worlds are not the result of free-for-all combinatorics – is a close variant of the *metaphysical* arguments found in Sellars' *Concepts As Involving Laws and Indispensable Without Them*.

²² For more general considerations on the relation between meaning and cognitive significance, see the appendix to §2.3.1 below.

²³ For a prominent example of this, see Soames' *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*.

§1 Possible Worlds and Semantic Modeling

The rigidity thesis for proper names by itself needs not decide the issue. For one might acknowledge that a proper name refers to the same thing (supposing that it refers at all) in all counterfactual situations, yet still believe that the modal behavior of a proper name is not all there is to its meaning. What does the doubt-casting is a group of so-called *epistemological* arguments.²⁴ One way to put the gist of these arguments is that, a proper name cannot be *synonymous* with a definite description whose content purportedly captures the cognitive significance of that name to a speaker (some speakers). For, the referent of the name is not determined *a priori* by that definite description, in the fashion, “the D, whatever/whoever it in fact is”. From Hitler to Columbus, to Gödel, Kripke argues, there is no description ‘D’ such that we cannot conceive the possibility that future empirical research reveals that the individual referred to by the proper name in each case in fact does not fall under ‘D’.

One place where this issue is hotly debated is the part of semantic literature dealing with the doctrine known as *two-dimensional semantic theories*, a particularly popularly pursued project amongst semantic theorists of late that attempts to reconcile (a) the Fregean idea that cognitive significance is part of meaning, and (b) the Kripkean lesson that dependence of truth and reference on ways the world might be needs not reveal differences in cognitive value (as the case of proper names demonstrate). These theories hope to accommodate both (a) and (b) by enriching the intensional assignment scheme with a second dimension of variation, one of *contexts*, in addition to that of possible worlds. The approach is inspired by Kaplan’s successful modeling of indexical meanings and his theory of “characters”.²⁵ Here is not the place to firmly settle the score one way or the other. But it is possible to get some idea of the complexity of the issue by looking at one version of such an attempt and reasons why it might not succeed.

As a project to model the cognitive significance component of meaning, while accommodating the Kripkean thesis of modal rigidity for proper names, a two dimensional semantic theory attempts to model that component as the dependence of reference for a proper name on contexts. The reference of a proper name will be fixed once the context is fixed, and will not further depend on possible worlds. Furthermore, the account may incorporate an insight of Gareth Evans so as to avoid colliding with Kripke’s epistemological arguments. Evans suggests that Kripke’s anti-descriptivist arguments need not be understood as precluding descriptions’ playing *any* role at all in the determination of reference. Rather,

²⁴ For the arguments themselves, see. pg. 71-97, Lecture II of *Naming and Necessity*, especially Kripke’s discussion of “Thesis (2)” and “Thesis (5)”. For the branding of these arguments as “epistemological arguments”, see for example Forbes, *Languages of Modality* pg. 122-129, and Soames, *Beyond Rigidity*, pg. 21-22, or Soames, *Ambitious Two-Dimensionalism*.

²⁵ C.f. Kaplan’s seminal essay „Demonstratives“

Part I: Semantic Preludes

the force of those arguments is that the *way* descriptions can participate in the determination of reference is not by supplying a *descriptive content* that the object referred to has to *fit*, but by being the content of an *information state* to which the object referred to has to *stand in a suitable causal relation*. (c.f. Gareth Evans, *The Causal Theory of Names*). A two-dimensional semantic theorist can take up this thought in his model and thereby respect the intuitions expressed in Kripke's epistemological arguments, in addition to those of his modal arguments. So such a semantic theorist might describe the variation of contexts as variation of the causal-histories of the language user's various information states. Given any token of proper name in use, there is a dominant piece of information the speaker associates with her use of the name. Now any choice of context will determine the reference of that name, not by finding an object *correctly described* by the descriptive content of that dominant information, but by locating the causal origin of the information state having that dominant piece of information as content. In other words, the context determines a particular variant of causal histories of *all* information states of the speaker, while the name selects the relevant information state to consider the causal history of.

Modeling the contribution of cognitive significance to meaning this way does not demand that the information a speaker associates with the name is *true* of the object referred to. For the relation between the information and reference is, to put it in a rough and ready form, not one of *fitting*, but one of *causing*. Kripke's epistemological arguments are therefore respected. But, it is not clear that what is being modeled is still the *cognitive* significance of a proper name, if the relevant descriptions need not even be *true* of the object referred to. For, what constitutes difference in "cognitive significance" modeled this way will have nothing to do with differences in descriptive content, but will be constituted by the brute numerical distinction between two tokens of information states, which, by virtue of being two numerically distinct states, not by virtue of the different descriptive contents they have, may have different causal origins given difference in context. Moreover, change of cognitive significance – under any reasonable interpretation – needs not imply change of meaning. So consider one possible future development in which one associates with "Hesperus" appearances in a particular kind of books in which the planet is always mentioned as such, not with observability at twilight. That is, we are imagining that *our current* mode of presentation for "Hesperus", which will have been only historically important for fixing the word's referent, will be replaced by other modes of presentation as history marches on. It seems fair to say that the *cognitive significance* of the word will have changed. It is not clear however that we want to say, anytime such changes take place, that the meaning of "Hesperus" has been altered.

For at least these two reasons, then, it is not clear that incorporating Evans insight into a two-dimensionalist framework will help re-habilitate the Fregean idea that cognitive

significance is part of meaning (at least for proper names). Now I will close the discussion on this topic without coming to a definitive conclusion on the matter.

§1.3.3 Possible Worlds As Representations

So instead of “stipulations”, we might think that we *represent* possibilities with possible worlds. According to this conception then, possible worlds merely *reflect* modal realities, so are constrained by these. This avoids the absurdity of “saying so makes it so” with regard to modal facts. On the other hand, like all representations, possible worlds are *formulated by us*, and this fact preserves something from the “stipulative understanding”.

A better variant of this conception takes possible worlds to be representations that are *possibly* accurate, rather than representations *of* possibilities.

§1.3.3.1 Possible Representations, vs. Representations of Possibles

To explain what I mean, let me back up a bit in the history of philosophy. There is a familiar problem, or feature, about the mental, as opposed to the physical, whose first articulation is widely attributed to Franz Brentano.²⁶ The feature is that a mental state can have a certain directedness towards some (intentional) object without that object actually existing. There have been by and large two ways of explaining this directedness to an intentional object that may not exist. One kind of explanation is based on distinguishing different *modes of existence* for the object directed at, so that an intentional object can merely exist in one mode (*bestehen*) without existing in the other mode (*existieren*). Another kind of explanation, which Brentano himself took up in his later writings, is to say that the relation of “directed at” is not *really a relation* – for a real relation would require the existence of all the relata, including the intentional object that needs not exist. Rather, the “relation” of directed-at is a *quasi-relation* (Brentano himself describes it as “etwas Relativliches”). The second kind of explanation distinguishes different modes for a relation to occur between the subject and ordinary objects, if you will, rather than distinguishing different ways for the intentional object to exist.

²⁶ C.f. Franz Brentano, *Psychologie vom empirischen Standpunkt*, ed. Oskar Kraus, 2 vols. (Hamburg: Felix Meiner, 1955). The famous passage is located on I.124–5/88.

Part I: Semantic Preludes

I think the second kind of explanation is clearly the superior one. The gist of that explanation can be put both more generally and more precisely, as follows. Many mental states and linguistic acts have a *representational component*. For mental states this might be mental representations, for a linguistic act, this might be the utterance made. The representational component represents the world to be a certain way, and to correctly grasp the corresponding state or act, one needs to grasp which way it is that its representational component represents the world to be. But different mental states or linguistic acts have different *modes* or – a not coincidental cognate of that word – *moods*, which we can represent as separate, and having a scope over the representational component, as follows:

- (5.a) *X believes/claims* that R represents accurately;
- (5.b) *X takes/asserts as possible* that the R represents accurately;
- (5.c) *X fears* that R represents accurately;
- (5.d) *X hopes* that R represents accurately.

So there is not a binary relation between X and some intentional objects at all in any of these cases. There is a *property – a monadic relation –* of “representing accurately”. Now if that property in fact obtains, and if R contains representations of objects and events that do not in fact exist, then indeed, we would have to entertain *binary relations* between R, or the subject, on the one hand, and the objects and events represented by R, that however do not exist, on the other hand, thereby landing in a paradox. But the property of accurately representing could not in fact obtain, if the objects and events represented by R do not in fact exist. More generally, the fact that the property of accurate representation is predicated within the scope of the “mode-verbs” (believes/claims/takes as possible/fears/hopes etc.) means that the truth of any of these four sentences does not imply that R in fact has the property of truly representing. So there is no need to entertain binary relations and intentional objects at all.

Now some (e.g. Christopher Menzel, see below) have proposed that we understand possible worlds in the same way. So to say that there is some possible world in which D, where ‘D’ is a potentially long and complex description, is to say that:

- (6) It is possible that ‘D’ represents accurately.**

This is to be contrasted with:

(7) ‘D’ represents a possibility.

where ‘possible’ or its cognate occurs within the scope of ‘represents’. The absurdity of understanding the possible world claim – that there is a possible world in which D – this way, literally, can be brought out by the parallels to the sentences in (5):

(8.a) X represents with R a *believed object*;

(8.b) X represents with R an *entertained-as-possible object*;

(8.c) X represents with R a *feared object*;

(8.d) X represents with R a *hoped-for object*.

It is true that (7) differs from (8) in that (7), literally interpreted, requires a *modal* mode of existence, whereas (8) requires, let us say, *attitudinal* modes of existence, both as *alternatives* to *actual* existence (the way “bestehen” is supposed to be an alternative to “existieren”). But the absurdity is of the same kind. (My argument does not however exclude the possibility of understanding (7) *non-literally*. See below.)

A particularly clear representative in the literature for the view represented by sentence (6) is Christopher Menzel, who writes in his essay *Actualism, Ontological Commitment, and Possible World Semantics*:

My family situation, for example, could have been much as depicted above, except that in addition I could have had a third child. Once again, I could draw a picture like the one above, this time with five figures rather than only four, the fifth labeled with, say, an X. Now in what way does this represent the possibility at hand? Not, I want to insist, in virtue of the fifth figure representing some merely possible family member the way the other four represent actual members. Rather, it is in virtue of a modal fact about the entire drawing: the drawing *could have been* such as to represent how things stand with my family. i.e., it could have been the case that each figure in the drawing represented some family member, and that each family member was represented by some figure in the drawing...

– p.371

Menzel is explaining what it is to claim that a picture represents a possibility, specifically, the possibility that his family included, besides his wife, his two children, and himself, a third child. The point is supposed to be about not only pictures, but also

Part I: Semantic Preludes

models that represent possibilities. His presentation is not explicitly formal regarding this point, but one can crystallize it by saying that the *correct analysis of the logical form* for the sentence “R represents a possibility”, where R is either a picture or a model, is *not*

Exist P (Possibility(P) & **Represent** (R, P));

with **Represent** being a binary relation, but *rather*:

Possibly (**Represent** (R))

with **Represent** being a *monadic* relation and having the meaning of *represent accurately*, and **Possibly** a *sentential* operator.

§1.3.4 Possible Worlds as Heuristics For Semantic Modeling

Now the choice to use representations that are possibly accurate, as described above, has much to speak for it. But what speaks for it is *not* that they capture the “nature” of possible worlds well. The “possible worlds”, as we recall, is a *place-holder* for whatever we *decide* to use to model meaning. The process of choosing a specific way to fill up the place-holder is not like the process of investigating the “nature” of possible worlds in a way that is similar to chemists investigating a particular compound, or the way legal scholars examine the content and implication of a specific clause in the American constitution. What we do is to decide how we want our assignment scheme to look in detail, in order that it *models meaning better* than our previous, extensional assignment scheme. Since the issue about possible worlds is embedded in a larger project of *modeling* meaning, the question must be: what available material do we choose to use and how we choose to use it for modeling meaning? Our task *cannot* be one of trying to reveal the true-nature of a bit of material used to do the modeling. Establishing the correct methodological attitude towards possible worlds is essential for avoiding the unnecessary metaphysical tussles later on.

But what *is* the correct methodological attitude? What guide do we have in making our choice, apart from the ability of the resulting assignment scheme to better model meaning? The place-holder “possible worlds” is introduced, as we saw, as a strategy to improve *one particular aspect* of the modeling. The strategy takes note of the fact that truth and reference depend on not only meaning, but also on possible

ways the world might be. It suggests assigning not the *actual* extension to words, but extensions for different ways the world might be.

But the strategy can be thought of as a mere heuristic guide for rectifying concrete modeling failures of the extensional scheme: the failure to register intuitively felt meaning differences, and the failure to model entailment relationships. In other words, though we can use *anything* to fill the place-holder “possible worlds” (and some wild choices *have* been suggested), to *call* what we choose “possible worlds” is to *let ourselves be guided by certain heuristic*, that, when we follow it to construct the assignment schemes, makes success of modeling more likely.

§1.3.4.1 Epistemic vs. Explanatory Priority

Against the view that the name “possible worlds” is only a heuristic, one might think that the way we motivated the introduction of possible worlds crucially makes use of a certain *priority* of possible worlds over meaning facts. For it is said that we *show or demonstrate* the fact that “sleep” and “snort” are different in meaning by pointing out that it is possible for the world to be in such a way that someone is sleeping but not snorting. This impression seems to fly in the face of the methodological attitude just suggested. It also flies in the face of one of the basic premises of this essay, namely, that possible worlds cannot serve as the basis upon which either meaning facts or modal facts are explained.

But there is in fact no such conflict. The priority of “ways the world might be” over facts of meaning difference can be thought of as an *epistemic priority*, rather than an explanatory one. Let me explain. The felt priority of possible worlds consists in the fact that, sometimes, in order to *know* or to *convince* ourselves and others of a certain meaning difference, we inquire about ways the world might be. If, let us suppose for a moment, An inferentialist such as Brandom is right and meaning does consist in inferential articulations of norms of linguistic use, *then*, in order to have knowledge about meaning differences, it might simply be *too hard* for a speaker to inquire about differences in inferential articulations of norms. If Brandom’s explanation of meaning is right, in other words, then it is simply much *easier* to judge about the possible ways the world might be, rather than directly inquire meaning facts, as Brandom describes it.

Judgment about possible ways the world might be is not *about* linguistic norms, but it involves the *actual exercise* of one’s grasp of these linguistic norms.

Part I: Semantic Preludes

And this explains why it is easier to inquire about possible ways the world might be, than to inquire about linguistic norms. The point can be illustrated with an analogy. If you ask a skilled unicyclist which position of the hands enables one to ride most stably, or a professional swimmer at which point one should turn one's head to breathe, the unicyclist or the professional swimmer would try to find out the answer to your question by *actually making the moves* herself. She could also examine videos of herself riding the unicycle or swimming. But that is a less efficient way to find out. Even less likely will she go to a drawing board and calculate the physics of it all. Now making judgments about ways the world might be involves actually exercising one's grasp of linguistic norm, for it takes knowledge of linguistic use to know that it is *not* possible that something is both a cat and a dog, that it is *not* possible that there are square triangles, and that it *is* possible for something to be sleeping but not snorting. (According to Brandom, to make these modal judgments is, in an important sense to be spelled out later in §3.4.2, *presupposes* a grasp of the inferential proprieties that supposedly constitute meaning.) Precisely *because* instances of such judging are exercises of one's linguistic understanding, the *correct* judgment on possible worlds "tracks" the norm of linguistic use of the expressions involved.

The analogy with the unicyclist and the swimmer is not exact because an exercise of the capacity for something (correct breathing time, norms of linguistic use), of which one desires knowledge, is, in the case of linguistic meaning, also an act having the *form of a representation*, that is, a *claim* that purports to be about possible ways the world might be. It is because of this that the question can arise at all about the priority of what this exercise appears to represent (possible worlds) over the original object about which one seeks knowledge (facts about meaning differences). Now we see that this priority need not be of explanatory nature. In fact, I have just suggested an explanation scheme in the *reverse* direction, *of* judgments about the ways of the world might be *in terms of* linguistic norms of use, which are, according to some theories, constitutive of meaning facts. The suggestion was: in judging certain ways to be possible or impossible we exercise our grasp of the proprieties of linguistic use, and consequently, the correct judgment will reflect these proprieties. In short then, the epistemic priority and explanatory priority come apart in the case of ways the world might be and meaning facts.

The proper way to appreciate the epistemic priority of possibilities over facts about meaning differences is to see that this is precisely what makes “possible worlds” such a *good heuristic*! Talks about possibilities come so naturally to us that the phrase “possible worlds” furnishes us a quick access to a rich panoply of pictures and thought patterns to help us make formal constructions of a modeling scheme.

§1.3.5 Which Sense of “Possible”?

Even at the level of heuristics, we can differentiate, on the one hand, the kind of possible-worlds employed in an intensional semantic modeling the way I have motivated it, and, on the other hand, possible worlds motivated by another kind of modeling need. In thinking of possible worlds – in the kind of intensional semantic theory discussed here – as “possible worlds”, we are thinking about a particular *kind* of possibility. The sense of “possible” we are thinking about is not “possible given the traffic laws”, or “possible given the energy resources available to men”, or “possible given our moral feelings”. In fact, if we want our judgments about ways the world might be to track meaning facts as much as possible, we should minimize constraints on these judgments over and above proper linguistic usage. For any admixture of these extra constraints, coming from legal systems, particular facts about the world, moral feelings, or what have you, will make our judgments not track meaning facts *alone*, but meaning facts *plus* these other given facts, whatever these are. So ideally, we want our judgment to be *only* constrained by proper usage of the linguistic expressions with which we entertain these possibilities, or in short, by the meaning of these expressions. Speaking loosely, we can say that the kind of “possible” is “possible, given (only) the meaning of the words used to specify these possibilities”. We might call this the notion of *conceptual possibilities*, where “conceptual” registers the fact that nothing is fixed except the *concepts* expressed by the words we use to articulate the possibilities.²⁷

²⁷ Now obviously, this notion of “conceptual possibilities” depends one’s theory of *meaning* – a theory that is explanatory. So if one one’s theory of meaning allows for meanings facts that are not knowable *a priori*, then the corresponding conceptual possibilities would also not entail *a priori* knowledge.

§1.4 Possible Worlds II – Modal Vocabulary

§1.4.1 Non-Extensional Contexts Generating Operators

The fact that extensional assignment schemes do not distinguish co-extensional expressions that differ in meaning was not identified as problematic due to a failure to model the *compositionality* of meaning. It was thought as problematic because it led to a failure to model entailment and intuitions about meaning differences. But if the object language contains a sentential operator that creates a so-called “non-extensional context”, even the modeling of compositionality will be jeopardized. A non-extensional context is a context in which expressions whose extensional features (reference and truth) are exactly the same can combine with the operator to yield compound sentences receiving different assignments. A special class of such contexts are also characterized, after Quine, in terms of the failure of substitution *salve veritate*. Examples of such operators are ‘it is (conceptually) possible that’, ‘Maria believes that’, ‘it is nomologically necessary that’ etc. If the object language contains one of these expressions, no extensional assignment scheme can model the meaning facts of that language in a way that conforms to a principle of compositional modeling like the one we formulated earlier²⁸.

The presence of non-extensional context generating operators gives rise to another motive for introducing “possible worlds”. Consider first the operator “it is conceptually possible that”. The goal is to modify the extensional assignment system to make it again capable of modeling meaning *compositionally*. However the modification goes, it is reasonable to assume that the resulting system distinguishes at least the extensional differences between expressions, that is, it will still distinguish sentences having different truth-values, noun phrases having different references, and predicates having different extensions. The following two sentences differ famously in truth-value:

- (1) It is conceptually possible that the number of planets is less than 9.
- (2) It is conceptually possible that 9 is less than 9.

On the other hand, the sentences following the operator “It is conceptually possible” have the same truth-value:

²⁸ C.f. §1.2.2.1 above.

- (3) The number of planets is less than 9.
- (4) 9 is less than 9.

These facts, together with our decision to continue to let our assignment system distinguish extensional facts of expressions, mean that we need *at least* to change the assignment system in such a way that the two component sentences in (1) and (2) after the possibility operator are assigned *different* things. That is, the assignment to component sentences after the possibility operator should be *more discriminating* than the truth-values these sentences *would* receive in the old extensional assignment scheme, had they been *free-standing*.

This leaves us at least two different choices for upgrading of our assignment system. We can make our assignment to sentences be sensitive to their status as free-standing or occurring in non-extensional contexts, and give the latter sentential occurrences a more fine-grained assignment. That is, we can make sentential assignments *context dependent*. Alternatively, we could upgrade the assignment of *all sentences*, regardless of how or where they occur.

The first option, together with the concept of *indirect contexts*, goes back to Frege. There are two arguments in favor the latter, non-Fregean option. First, the Fregean option comes with the conceptual distinction of sense vs. reference, a theory of two-tiered semantic values, as well as a theory of how these are related in indirect contexts. But this theory is best motivated for contexts created by epistemic verbs, not modal contexts strictly so-called. Recall Kripke's point that epistemic contexts and modal contexts differ greatly in what they are sensitive to (as shown by the case of proper names). We saw how it gives rise to doubts as to the naturalness of the two-tiered semantic hierarchy (c.f. Forbes' exploitation of these facts to argue that "states of affairs", rather than truth-values, are the better choice for sentential "references" in the Fregean two-tiered scheme, in *Languages of Modality*, which we discussed briefly above). Second, regardless of the expressive power of the object language, extensional assignment to sentences falls short of distinguishing true sentences that are different in meaning. This we have seen in the previous section. So there is a motivation for refining sentential assignment *across the board*, even before we consider the task of accommodating object language richer in expressive power.

Either way, the assignment scheme will have to be refined for *some* sentences. In case the trouble-making operator is ‘it is conceptually possible that’, the refined assignment must be able to detect differences in truth-value *for any conceptual possibilities*. By a **principle of economy**²⁹, the refinement should be just fine enough to detect such differences. This leads again to *intensional* assignment schemes where, instead of truth-values or objects, we assign *functions* from conceptual possibilities to these. The situation is similar when the operator in question is *modal*, such as “It is permissible that”, “It is nomologically necessary that”, etc. For these modal operators, the possible worlds we invoke will not be thought of as *conceptual* possibilities. Instead, we arrive at an *intensional* assignment scheme in the sense that *functions* from possibilities are assigned. The case of operators that make attitude attributions has their special difficulties and will not be figure in my discussion here.³⁰ In summary then, when the object language has modal expressivity, our desire to model compositionality of meaning forces us to adopt an intensional assignment scheme.

§1.4.1 Variety and Iteration

Two things make possible worlds as basis for modeling modal language different from possible worlds as basis for intensional modeling of non-modal language. The first is the fact that there are a variety of different modal operators, each having a different meaning. The second is the fact that at least some of them can be iterated.

§1.4.1a Variety of Operators

First, different non-extensional context generating operators clearly demand different intensional assignment schemes based on different possible worlds heuristics. So if the operator is “It is nomologically possible that”, the required heuristic of possible worlds, for the purpose of modeling the interaction of this operator with the rest of the object language compositionally, will be based on nomological possibilities, that is, the notion of what is possible *given the laws of nature*. Similarly, for the operator “It is legally permissible that”, the requisite heuristic of possible worlds will be based on legal permissibility, which we can gloss as the notion of what is possible *given the*

²⁹ Rather than the full-blown *context principle*, which is used to determine not only the correct assignment, but also the correct classification of grammatical categories, such as referential vs. non-referential.

³⁰ See the footnote 32 below for more details on the special problems attending these expressions as well as Hintikka semantics.

laws (in the sense of the content of legal documents). And for the operator with which we began this section, “It is conceptually possible”, we can gloss the notion of possibility required as the notion of what is possible *given semantic laws* (i.e. meaning facts). In general, then, for any one of a large family of modal operators, there will be a corresponding category of laws, call it **F**, so that, to correctly model the interaction of this operator with the non-modal fragment compositionally we need a “possible worlds” heuristic based on the notion of what is possible *given, or respecting, the F-laws*.

§1.4.1b Iteration of Operators

Second, the fact that a modal operator might be iterable means that the question “*which F-laws?*” is not an idle one. For it cannot always be the *actual F-laws* that are meant. Consider the sentence:

(S**) It is nomologically possible that it is nomologically possible that S.

The truth value of (S**) will be sensitive to, thanks to the first operator, the extensional facts (in this case, truth) about the embedded sentence

(S*) It is nomologically possible that S.

relative to all nomological possibilities, conceived as “possible under *actual* laws of nature”. Suppose P is such a nomological possibility, we will need to know if (S*) is true in P. This is something novel for us, because it means we have to allow for *modal truth in a possible world*, as S* has the modal form of “It is nomologically possible that...”. This appears to show that, at least for those modal operators that are iterable, possible worlds whose description fixes a relevant class of *modal* facts are needed to model their semantics.

Furthermore, the truth-value of S* in P, is the truth-value of S* given that the world had been P. That in turn is sensitive to the *truth-value of S* relative to all nomological possibilities, again, supposing the world had been P. Which nomological possibilities are these? Intuitively, the relevant ones are those understood as “nomologically possible, supposing that the world had been P”. We noted just now that by supposing the world to be P, we fix facts of what is nomologically possible and what is not. So the relevant nomological possibilities invoked in evaluating the truth-value of S in P must be those *fixed by P* to be nomologically possible, rather

Part I: Semantic Preludes

than what is *actually* nomologically possible. Considerations such as this leads to the introduction of an *accessibility relation* among modal possible worlds to model the idea that descriptions of possible worlds can contain modal claims themselves. there is therefore one more difference between the “possible worlds” heuristic for intensional modeling of non-modal languages and the “possible worlds” heuristic for modeling modal language.³¹

To mark these differences, especially the latter one due to iterability, let me call the possible worlds used for modeling a language with modal expressivity *modal possible worlds*, and possible worlds used for intensional semantic modeling of extensional languages *non-modal possible worlds*.

Given this terminology, we can say that, the *modal* conceptual possible worlds are, in themselves, as it were, the same as *non-modal* conceptual possible worlds (they are conceptual possibilities). But collectively, as *modeling device*, they differ in that modal conceptual possible worlds are *interconnected* by conceptual accessibility relations, while non-modal conceptual possible worlds are not (*even though* the relation in this case is the trivial relation that relates everything to everything else). Here we see again that it is inappropriate to ask the *nature* of possible worlds (“are they by nature related by accessibility relations?”). What they are solely depends on the modeling task we use them for.

Despite the differences, there is a sense that non-modal possible worlds are more basic than all the modal possible worlds. In fact, I will argue that the intensional assignment scheme we are *forced* to adopt for an object language with modal expressivity is a straightforward extension of the intensional assignment scheme we would adopt for its *non-modal* fragment motivated by accidental co-extensionality.

§1.4.3 Non-Modal Conceptual Possible Worlds Are Most Basic

Consider again what I have been calling “conceptual possibilities” as the basis for non-modal possible worlds. What is the difference between these and the *modal* possibilities as a basis for modeling the operator “It is conceptually possible that...”?

³¹ There is a distinction between *global* and *local* nomological possibilities, depending on whether we are talking about possibility *given the meaning* of “nomological”, or possibility *given a particular set of laws of nature* – which can intuitively be thought of as a *referent* of “nomological”. The local nomological possibilities are relative to the nomological possible worlds, while the global nomological possibilities *delimit* the nomological possible worlds. For more on this distinction, see discussion in §1.4.3 below.

Under the current nomenclature, the latter are maximally specified – including conceptual modal facts – possibilities. Two things can be said about these modal possible worlds. First, the words used to specify them, are always used *with the same meaning*, no matter which modal possible world it is that is being specified, and *even when* these words occur within the scope of the modal operator. So in describing a modal (conceptual) possibility with the modal sentence “it is conceptually possible that someone is sleeping but not snorting” the words “sleeping” and “snorting” are meant to be used with their *normal* meaning.³² So the specification of modal possible worlds is subject to the same constraints as is the specification of *non-modal* conceptual possible worlds. This point in fact holds for *all* kinds of modal possible worlds. For surely, in saying that it is nomologically possible that black-holes will result from the experiments in CERN and slowly devour the entire earth, we do not mean to use the words “black-hole”, “earth”, or “devour” with some queer, nonstandard meanings. And that is so despite the fact that the *possibility* we want to describe is a strange and extra-ordinary one. We mean to talk about *black-holes*, normally understood, *devouring*, in ways “devour” is normally understood, *earth*, again, as the word is normally understood. We certainly do not mean to say with these words that it is nomologically possible that cows will migrate to the Antarctic, reading “cows” for “black-holes”, “migrate” for “devour”, and “the Antarctic” for “the earth”. The same goes for any other kind of possibilities. In specifying an alternative legal possibility, we intend our words to mean the same as when we use them to describe the actual situation, for example.

Second, the converse is also true in the case of *conceptual* modal possible worlds. That is, specification of conceptual modal possible worlds is *only* subject to the sameness of meaning constraint (i.e. same from world to world). It is true that descriptions for modal conceptual possible worlds are richer in that they contain conceptual modal claims. But these conceptual modal claims are just claims that *make explicit* the proprieties of use of the meta-language, with which the specifications of worlds is made. So they *have* to be true also in *non-modal* worlds. It is just that non-

³² Note that this argument – that the world specification uses expressions conforming to actual semantic rules – cannot be made had the modal operator been an *attitude* report operator. This is the chief reason why modeling the meaning of attitude sentences are so much harder. A Hintikka style semantic theory cannot be based on what I here call “conceptual possible worlds”. It is unclear *what* conception of possible worlds, if any, should be adopted for that purpose. I avoid discussion of this matter altogether in this essay.

Part I: Semantic Preludes

modal worlds, to serve their modeling purpose, need not *explicitly contain* these claims.

The upshot of these considerations can be put as follows. Consider three languages: $\mathbf{O} \subseteq \mathbf{OC} \subseteq \mathbf{H}$, where \mathbf{OC} is the object language with “It is conceptually possible that...”; \mathbf{O} is the non-modal fragment of \mathbf{OC} , and \mathbf{H} is the meta-language. The non-modal possible worlds used for intensional modeling of \mathbf{O} can be 1-to-1 mapped onto the modal possible worlds used for modeling the modal language \mathbf{OC} . The only difference between them is that the description of each modal conceptual possible world contains an identical set of conceptual modal claims, and that every modal conceptual world is related to every other by accessibility relations.

In case the modality in question is not conceptual, the transition from non-modal possible worlds to modal ones will be non-trivial, and involve *differentiation*, and, depending on the modality, possibly also *selection*. So the description of a non-modal possible world may be enriched with modal facts in more than one ways (e.g. different laws of nature can fit the same set of non-modal data gathered), in which case one non-modal world is differentiated into several nomological worlds. Moreover, it may happen, *prima facie* at least, that some non-modal conceptual possibility does not admit any “nomologization”. In that case, that non-modal world will not pass the “selection” to become a nomological world. In any case, given any modality, there is a natural map from the set of modal possible worlds to the set of pre-modal possible worlds, via the process of “forgetting the modal facts”.

In fact, the precise manner of differentiation of and selection from the non-modal conceptual worlds (the accessibility relation follows from these) models the meaning of the respective *modal operator*. So the *modal* possible worlds are also *conceptual*, the only difference being that (a) one of the concepts fixed here is a *modal concept*, namely, the concept expressed by the modal operator in the object language in question; and (b) modal possible worlds are specified with more detailed specification (extra *modal* specification). For each modal concept then, there is a corresponding notion of modal possibilities or modal possible worlds, fixed by normal conceptual possibilities plus the meaning of that modal concept.

In case the modal concept in question is not the simple conceptual modality, we need to carefully distinguish between this and a previously mentioned notion of possibility. Again, let nomological modality serve as an example. The notion

previously mentioned is what we have been describing as “possible given laws of nature”, and it is a notion *relative* to a particular nomological possible world, because the reference of “laws of nature” are, *prima facie*, relative to them. The notion of conceptual possibility, holding fixed the concept of nomological modality, is “possible, holding fixed the meaning of words used in our description, including the modal word ‘nomological’”. This *latter* notion of nomological possibility is clearly *absolute*, that is, *independent* of any nomological possible world. In this sense, all nomological possible worlds, whether or not they are “accessible” from the world that describes the way the world *actually* is, are nomologically possible. Let me distinguish these two notions by calling the former *local nomological possibility*, and the latter *global nomological possibility*. When I am omitting the local vs. global specification, I mean to refer to the *global* notion.

§1.5 Conclusions

Possible worlds, as we saw in the above sketches of model-theoretic semantic theories, naturally emerge as a basis for constructing a systematic compositional modeling of the meaning facts of a language, in the assignment-cum-functional application style. They are invoked first to deal with accidental co-extensionality through an intensional refinement of our assignment scheme, regardless of whether the language we seek to model can express modal notions. They are invoked a second time, and this time practically forced upon us because of our commitment to compositionality, when we expressly desire to model the meaning of language capable of expressing modal facts.

The way possible worlds naturally suggests itself as a basis for semantic modeling does not warrant making postulating materially parallel existing worlds. In fact, they can be taken to be representations of aspects of the world that are *possibly accurate*, and complete in some sense. But to fill the place-holder “possible worlds” this way is not to make claims on the nature of a special kind of entities. Other ways of filling the place-holder is conceivable. Calling possible worlds “possible worlds” is to adopt a useful heuristic device for constructing successful modeling schemes.

The simple heuristic requires a prior notion of conceptual possibilities. For some modeling purposes – say for a modal object language – more refined heuristics are needed, presupposing accordingly different notions of possibilities. The actual

Part I: Semantic Preludes

formal device for modeling becomes more complex (for example in having accessibility relations due to the need to model iterability of modal operators). But because of the interconnections among the corresponding conceptions of possibility, the heuristics for more complex modeling than be thought of as derived from the heuristic of conceptual possible worlds, through the process of differentiation and selection.

In any case though, because of the employment of “possible worlds” heuristic is to assist semantic modeling, rather than to provide semantic explanations of modal meanings, its application involves no vicious circularity and the like. The rest of the essay, insofar as it deals in semantics, will be devoted to the *explanatory* questions. It begins by surveying alternative semantic tradition for ideas.

§2 Semantic Functionalism – A Selective Overview

§2.1 Introduction: Two Sources and Four Choices	46
§2.2 Two Sources – Dummett and Sellars	48
§2.2.1 Dummett: Metaphysics vs. Occultness.....	51
§2.2.1.a Dummett & Wittgenstein – The Turn To Understanding.....	51
§2.2.1.a(i) Dummett vs. Ryle.....	52
§2.2.1.a(ii) Epistemology of Meanings.....	53
§2.2.1.a(iii) Manifestation of Linguistic Understanding	53
§2.2.1.a(iv) Verificationism	54
§2.2.1.a(v) Argument Against Truth-Conditional Semantics.....	55
§2.2.1.a(vi) Late Wittgenstein	57
§2.2.1.b Priority of Understanding.....	59
§2.2.1.b(i) Dummett in His Own Words.....	59
§2.2.1.b(ii) Fodor’s Ridicule	62
§2.2.1.b(iii) More on Fodor’s Misunderstanding of Dummett.....	64
§2.2.1.b(iv) Arguments For The Priority-Of-Understanding Thesis.....	66
Appendix to §2.2.1.b: Meaning, Understanding, and Communication.....	69
§2.2.2 Sellars: Meaning Nominalism	71
§2.2.2.a Meaning Nominalism and Functions.....	72
§2.2.2.b Functions and Their Pragmatist Interpretation.....	75
§2.2.2.c The Order of Being vs. The Order of Knowing	78
§2.3 Content of Thoughts vs. Meaning of Expressions	80
§2.3.1 Varieties of Positions.....	80
Appendix to §2.3.1.....	81
Meaning, Content, and Cognitive Significance.....	83
§2.3.2 Taking Stock.....	84
§2.4 Which Semantic Functionalism?	85
§2.4.1 Two Kinds of Semantic Explanations	85
§2.4.2 Dummett – Implicit Knowledge/Consciousness	89
§2.4.3 Sellars – Experience & “Conception Of The Norm”	91
§2.4.4 Brandom – Mindless Meaningfulness	96
§2.4.4.1 Functional Description & Semantic Reductionism	96
§2.4.4.2 Mindless Meaningfulness.....	100
§2.5 Conclusions.....	104

§2.1 Introduction: Two Sources and Four Choices

An alternative approach to linguistic meaning based on the notion of inference rather than the notion of truth or truth-conditions has been developed in the second half of the twentieth century. It is this tradition that I shall extract the main ideas from for an account of nomological modal meaning. I will use the term “semantic functionalism” as a generic term to cover the family of theories developed within this approach, though as we will soon see, these theories can differ widely on motivation and explanatory objectives, and some are not motivated by functionalism at all. Some people³³ trace this general approach to linguistic meaning to the publication of *Philosophical Investigations* by Wittgenstein in the early fifties. But the earliest explicit defenders of inference as the basis for semantic explanation were Wilfrid Sellars and Michael Dummett. There are other notable connections besides the writings of late Wittgenstein. Some varieties of inferential semantic theories were tied up from the very beginning with functionalism in the philosophy of mind. The first part of this chapter (§2.2) will be a brief survey of these two basic, rather distinct ideas that have pulsed through the development of inferential semantics: (a) the change of focus, under the influence of Wittgenstein’s later work, from “meanings” as entities to linguistic *understanding* on the one hand, and (b) the application of functionalist ideas to the theory of linguistic meaning on the other. The main protagonists of this survey are Dummett and Sellars, each of whom, as I shall argue, stands for one of these two lines of development.

This difference within inferential semantic theories, corresponding to late-Wittgenstein and functionalism in the philosophy of mind, may have been obscured by the fact that the most prominent recent proponent of inferential semantics, Robert Brandom, has, by way of an encyclopedic approach to philosophy rather typical of his writings since *Making It Explicit*, fused these two lines of development. Details will emerge at the end of our historical survey. But roughly, what happened is that Brandom *both*, in a move reminiscent of Dummett, makes inferential linguistic practice alone do the explanatory work for linguistic understanding, *and*, in a typical functionalist move, he identifies specific “meanings” with structural features of that

³³ Notably the chief defender and articulator of the inferential vision, Robert Brandom, claims this lineage in *Making It Explicit*.

practice.³⁴ In any event, the *second* goal of this chapter, next to trying to lay bare the historical beginnings, is to delineate more finely the semantic ambitions of this essay, above and beyond the description “functionalist semantic treatment of nomological modal meaning”. Dummett, Sellars, Brandom as well as other philosophers that will come up in the following discussion all serve, in this context, as instruments of delineation and demarcation.

There are *four* basic parameters with which the goal of this essay will be more narrowly defined. They concern the following aspects of a theory of meaning, respectively:

- (C1) meaning vs. grasping meaning;
- (C2) ontological status of “meanings”;
- (C3) content vs. meaning;
- (C4) identity vs. *ti-esti* questions about meaning.

The first parameter is about being primarily a theory about the grasp of linguistic meaning vs. being primarily a theory about more strictly meaning-facts, such as identity and difference of meanings as well as polysemy. For this essay, the semantic theorizing will be about more strictly meaning-facts. The second parameter is about various attitudes towards the ontological status of “meanings”: eliminate them, reduce them (to “inferential roles”), treat them as irreducible abstract entities, etc. In this essay I will remain agnostic about whether there are “meanings” and, if so, what they are. The third parameter is a choice between, on the one hand, theories that purport to be about a generic concept of “content”, including both meaning of linguistic expressions as well as representational content of mental occurrences, and, on the other hand, theories that are strictly about linguistic meaning. The semantic discussions of this essay will be strictly about linguistic meaning. The fourth parameter is about explanatory goal. The choice is primarily between I shall call the *identity* explanation³⁵ of meanings on the one hand, and an account of what linguistic

³⁴ Both moves are in my view untenable. I will however merely indicate my suspicions why they fail rather than developing my own complete theories of linguistic understanding and of “meanings”.

³⁵ With the phrase “identity explanation”, which I shall explain in detail below, I do *not* mean to refer to explanation of the *identity conditions* for meanings as entities. As we shall see below, identity explanations as I define it, is independent of doctrines about the ontological status of “meanings”.

Part I: Semantic Preludes

meaningfulness as such consists in, on the other. The explanation relevant for this essay falls under the former, more modest rubric, not the latter.

The first choice is really a cutting down on ambition: I agree with Dummett that the central task of a theory of meaning is to explain the grasp of meaning, but I will not attempt such an explanation here. Among other things, theories of linguistic understanding are intimately tied up with the philosophy of mind, which is simply not the topic here. The second choice is made on the ground that there is in our ordinary usage no explicit referential device for “meanings”³⁶, but that there can none-the-less be considerable theoretical utility in making systematic assignment of “semantic correlates” to expressions of a language (for the purpose, let us say, of modeling its inferential structure). The third choice is made on the ground that “content” and “meaning” are, *prima facie*, two distinct categories, and that there is simply no space for a wide-ranging treatment of both in a single essay of this size. I will, however, sketch arguments for the externalist character of the concept of meaning, which would distinguish it from the concept of content. The fourth choice is more or less decided by the first choice insofar as an account of the linguistic meaningfulness as such cannot skip over the issue of linguistic understanding.

The chapter is organized as follows. In §2.2, the two historical/systematic sources of semantic functionalism will be examined. In the course of this largely historical survey, the substantive issues of (C1) and (C2) of the four parameters will also be touched upon. §2.3 is about the contrast (C3), between content of thoughts and meanings of words. The issue of linguistic understanding (parameter (C1)) returns in §2.4, and it is discussed against the background of distinction (C4), between the identity and *ti-esti* semantic explanations. The latter distinction has strategic significance, for it allows us to define a kind of semantic explanation largely free of the nexus of complex issues about consciousness, mind, and linguistic understanding.

§2.2 Two Sources – Dummett and Sellars

The first of the two motivations for semantic functionalism is a *metaphysical* one and it is rooted in a discomfort about the ontological status of the supposedly abstract entities “meanings”. The idea begins with the thought that the functionalist strategy

³⁶ E.g., we say “E means ...”, not “The meaning of E is (=)...”

for dealing with mental items, for whose ontological status there was equally a discomfort, might be fruitfully applied to the theory of meaning also.

The strategy in question is therefore a variety of *functionalism*. The conditions for applying functionalism are, generally speaking, as follows. In the beginning one is faced with two sets of phenomena, let us call them the *target phenomena T* and the *fundamental domain F* respectively. As a basic reference case we can think of T as the mental phenomena, and F as the phenomena described by physics or perhaps neurology. It is thought that there is an *ontological dependence* of T on F, perhaps some kind of supervenience. The idea is that, if items in F are all fixed, then so are items in T. But still, there is no obvious strategy to *re-formulate* descriptions of items in T into some complex descriptions for items in F. Nor does one want to eliminate the target phenomena, by claiming that our beliefs about T-phenomena, though serving some practical purpose, is either non-sense or systematically false.

Given a situation like this, functionalism is an attractive option because, if successful, it allows us to have a theory that (i) reckons the irreducible and ineliminable reality of the target phenomena T and yet (ii) respects the intuition that items in F are really *all there is*, that items in T are not, nor do they presuppose, *existences in addition to those in F*. Specifically, the functionalist strategy aims (a) to describe some sort of *functional system* made up of items in the fundamental domain and (b) to identify items of target phenomena simply as items from the fundamental domain *playing certain functional roles*. So in classical functionalism in the philosophy of mind it is claimed that a belief or a desire is that (neurological state) which plays a certain causal role in a complex causal system, part of which are outward behavior and items in the subject's immediate environment. The application to semantics will have the general form: for a symbol-token/psychological state (depending on one's view about the primary content/meaning bearers) to have representational content/meaning p is for that token/state to have such and such a role in a functional system.

We will turn to a more detailed exposition of functionalism applied to the theory of meaning in a moment.³⁷ For now, let us observe that the first motivation for inferential semantics arises from (a) a certain metaphysical view about what is

³⁷ See especially §2.2.2.b Functions and Their Pragmatist Interpretation below for the attractions of the functionalist strategy.

Part I: Semantic Preludes

ontologically more basic and what is less basic, combined with (b) a recognition that the less basic items cannot be descriptively reduced to the more basic ones nor be eliminated altogether.³⁸

The second motivation for inferential semantics is, by contrast, not an entirely metaphysical one, and it turns on a supposed feature of linguistic meaning that has been forcefully emphasized by Dummett. The idea is described and defended in various writings by Dummett, but the following short summary suffices for now:

...a theory of meaning is a theory of understanding; that is, what a theory of meaning has to give an account of is what it is that someone knows when he knows the language, that is, when he knows the meanings of the expressions and sentences of the language.

– Dummett, *What Is a Theory of Meaning (I)*, p.3

At issue is the supposed relative *explanatory priority* of a person's *understanding* of meaningful sentences/statements over meanings as such. Let us call that the Priority of Understanding Thesis about linguistic meaning. This, coupled with what I shall call the Manifestation Thesis about understanding, leads very quickly to the view that linguistic meaning should be explained by describing what one has to be able to do with a given symbol in order to count as understanding its meaning. The Manifestation Thesis is the view that the *understanding* of meaning or meaningful items (symbols) must be *manifested* in one's practical ability to do things with these items. The Priority of Understanding Thesis says where the central explanatory task lies for a theory of meaning, while the Manifestation Thesis specifies the *form* any such explanation has to take in order to be successful (i.e. a successful explanation must have the form of providing a description of how linguistic understanding is manifested).

There are many differences between these two ways of motivating inferential semantics. Some of them are: functionalist motivation is metaphysical whereas Dummett's motivation has an important epistemological component; the functionalist motivation tends to lead to theories about *contents* of mental items as well as theories

³⁸ As we shall see below, Fodor shares with the inferential role semanticists the *metaphysical judgment* about the fundamentality of physics and/or neurology. But he differs from them in his pursuit of a *reductionist* strategy for content and meaning.

of meaning whereas Dummett's can only be interpreted semantically; a functionalist theory is neutral on the centrality of *normativity* of linguistic meaning whereas Dummett's approach has a built-in commitment to it. The rest of the introduction will elaborate some of these differences, and along the way, we shall get a sharper focus on various prototypes of inferential semantic theories.

§2.2.1 Dummett: Metaphysics vs. Occultness

§2.2.1.a Dummett & Wittgenstein – *The Turn To Understanding*

The first and most important difference is that, whereas the functionalist motivation is based on the perceived problematic *metaphysical* status of meanings, Dummett's motivation, part of which he traces back to Wittgenstein's later writings, is centered around an attempt to replace the so-called "code"-conception of how language works, according to which a language's function is to encode language-independent thoughts. The conception that replaces it requires the workings of a language be completely "open to view". Dummett describes the code-conception of language thus:

Philosophers before Frege assumed ... that what a speaker knows is a kind of code. Concepts are coded into words and thoughts, which are compounded out of concepts, into sentences, whose structures mirror, by and large, the complexity of thoughts. We need language, on this view, only because we happen to lack the faculty of telepathy, that is, of the direct transmission of thoughts.

– Dummett, "*What do I Know When I Know a Language?*" p.97

The problem with this code-conception, according to Dummett, is that it relies on a metaphor that simply does not work. The metaphor is implicit in the talk of "coding", and the basis of the metaphor is the idea that a speaker associates bits of one language with bits of another language. But this "translation manual" picture ceases to make sense when one of the relata of "association" is meaning (or concepts/thoughts), rather than bits of language or other kinds of representations. For, unlike "a code" in "associate *a code* with an expression", the phrase "a meaning" in "associate *a meaning* with an expression" is syncategorematic. In other words, "a meaning" does not denote a bit of representation (e.g. a scribble of letters, a mental episode, etc.) that can stand in the same sort of relation to an expression the way a code can stand to it. Short of a specific account of what attaching a meaning to an expression consists in,

Part I: Semantic Preludes

one that can be put in place of the non-working metaphor of “coding”, no account of linguistic understanding is delivered. Moreover, the picture of linguistic communication entailed by the code-conception of language makes it a matter of *hypothesis* that my interlocutor attaches/associates the same meaning to an expression as I do.

One might say, therefore, that Dummett’s motivation depends on finding the idea of encoding thoughts – that language independent thoughts are encoded in linguistic expressions by individuals having these thoughts – an *occult* one. Since occultness is an *epistemological* notion, this seems to lead straight away to the contrast with which I have titled this section: namely, that the central semantic question for Dummett is not ontological status, but epistemological accessibility. This would be too simplistic and misleading, however, unless we add the following caveats.

§2.2.1.a(i) Dummett vs. Ryle

It might be thought that the occultness of the code-conception of language has two aspects: the occultness of *encoded thoughts*, and the occultness of *thought encoding*. These are, respectively, the occultness of *mental episodes* (thoughts) and the occultness of *understanding* (understanding as attaching thoughts to linguistic codes). Dummett’s semantic writings are focused on the second sort of occultness, rather than the first. A helpful contrast to Dummett is Gilbert Ryle, who, with his book *The Concept of Mind*, began a relentless attack on the occultness of mental episodes. Ryle’s key arguments against mental episodes make reference to their supposed role for giving actions an “intelligent character”. On some level, both Ryle and Dummett were attacking a certain traditional appeal to mental episodes to explain intelligent behavior. But this appeal has different roles in their respective theorizings. For Ryle, the appeal represents an argument from intelligent behavior to occult mental entities. Ryle’s objection to this argument derives from his objection to its conclusion: the existence of private mental entities. Dummett, on the other hand, objects to the appeal primarily because he thinks that, in the context of linguistic behavior at least, it *fails to explain what it purports to explain*, namely, linguistic understanding. This is also

why Dummett does not appear to reject the notion of mental episodes as such.³⁹ He rejects, more precisely, the appeal to them for certain explanatory purposes.

§2.2.1.a(ii) Epistemology of Meanings

In any event, what Dummett is determined to overcome is not the occultness of *meanings* per se, such as, to name another Titan in the tradition of analytical philosophy, Quine was. In fact, philosophers have found the ontological status of meanings troublesome precisely on the ground that there are purportedly no meaning facts one can come to know by the kind of *observation* that gives us knowledge of the physical world, i.e., the way one can observe the color of a plant or the hardness of a table. This supposed lack of epistemic access turns into a reason for metaphysical suspicion against meanings, as long as some empiricist principle or other is in play in one's philosophical reflections. Quine, for example, can be read along these lines: having assumed that the only observable items are behavioral and non-behavioral facts describable by natural science, Quine concludes that the only observable facts having a bearing on meaning are the direct causal correlations between proximate stimuli from the environment and instances of utterance-types caused by them. Accordingly, the only meaning-facts Quine is ready to countenance are what he calls stimulus-meanings.⁴⁰ Quine's idiosyncrasies aside, the general point is that metaphysical suspicion about meanings can very often be based on alleged epistemic opacity of meanings plus some version of empiricist principles. But that is not Dummett's concern.

§2.2.1.a(iii) Manifestation of Linguistic Understanding

Two things distinguish Dummett's epistemological concern from the generally empiricist one, corresponding to his Priority of Understanding Thesis and the Manifestation Thesis mentioned above. The first distinguishing character is that

³⁹ See, for example, his appeal to "reflections" discussed later in §2.4.2 Dummett – Implicit Knowledge/Consciousness.

⁴⁰ These two aspects of Quine's work in semantics – a general empiricist orientation and an extreme version of behaviorism that strictly limits what counts as observable – are in full display in his *opus magnum* "Word and Object". The work is framed methodologically by the *epistemological* question about meaning: on the basis of what evidence can the field linguist/radical translator come to what conclusions about the meaning of expressions in some native's language. In contrast, Dummett rejects the very talk of *evidence* for someone's attaching this or that meaning to an expression.

Part I: Semantic Preludes

Dummett is specifically concerned with the occultness of one's *grasp* of meanings – which according to Dummett is a consequence of the code-conception of language – rather than with the occultness of meanings in themselves. The second distinguishing character is that, unlike a (reasonable) empiricist who will be satisfied as long as he finds a plausible account of some *evidential* route to the disputed items in question, Dummett demands that the grasp of meaning, which is the item he is concerned with, be “open to view”. What this means is that Dummett's concern is not merely epistemological. For he would *not* be satisfied if one offered him an account of what in general counts as adequate *evidence* for concluding, inferentially as it were, that a speaker has grasped the meaning of a word. One of Dummett's complaints against the code-conception of language is precisely that, if it were accurate, it would remain a mere *hypothesis* that someone else understands the expression the same way as one does oneself, which is something needing (impossible) evidence for support. One might also say that Dummett demands an account of what counts as linguistic *understanding*, rather than what counts as *evidence* for such understanding.⁴¹ Dummett's own formulation is that a theory of meaning needs to specify how knowledge of linguistic meaning is *manifested in practical abilities*.

§2.2.1.a(iv) Verificationism

These two characteristics of Dummett's epistemological concern – having its locus on the grasp of meaning rather on meanings themselves *and* demanding not just a theory

⁴¹ McDowell insists famously that we *perceive* (hear) *directly* the meaning of a sentence addressed to us (see McDowell: *Criteria, Defeasibility, and Knowledge*). He may also have held – though he does not appear to have said so anywhere in print – that someone else's *knowledge* of meaning can also be *directly perceived*. If so, this would be a different position on the *epistemology* of linguistic understanding as Dummett's, with whom McDowell sparred about the correct form of a theory of meaning. For, facts that for Dummett *constitute* someone's linguistic understanding are rather complex, and are probably not directly perceivable in the moment of communication. The point being made in the main text is that these complex facts are *not merely evidence* for the fact that the person in question understands a certain expression, but *constitute* it.

The issue of “direct” perception concerns not only the *complexity* of the performance manifested, but also the problem of how to distinguish a performance “done with understanding” from one that is not. To say that understanding is manifested in *practical abilities* does not settle the matter really. For, as Dummett himself acknowledges, the abilities in question are not “purely practical”, such as the ability to ride a bicycle is, but involve the “awareness” of rules on the part of the speaker. (c.f. *What Do I Know When I Know a Language?* for Dummett's discussion on this point). To the extent that Dummett does not have an account of this “awareness” of rules that must accompany a linguistic performance done with understanding, he did not arrive at a definitive resolution of the Wittgensteinian question that informs much of his writings on theories of meaning. See below for a discussion of the bearing of this point to Dummett's ability to offer an explanation of the *ti-esti* of linguistic meaning.

of *evidence* for, but a theory of *manifestation/constitution* of understanding through practical abilities – should not only be contrasted with radical empiricist suspicion of meanings, they should also be contrasted with a doctrine for which Dummett himself is well known: verificationism. Without getting into any details of a historic or a systematic discussion on the matter, we can think of verificationism as empiricism transposed into the linguistic key. Verificationism concerns the meaningfulness, truth, or some other semantic character of linguistic expressions, whereas empiricism concerns the existence of items purportedly *meant/referred* to by a certain class of expressions. The former demands, for meaningfulness/truth (or some other determinate semantic status), that the truth-value of a given sentence be in principle epistemically accessible; the latter demands, for the purported existence of a kind of entities, that we have, in principle, epistemically significant interactions with – indeed “experiences” of – this type of entities. The possibility of *verifying the truth-value* of a sentence plays for verificationism the role played, in empiricism, by the possibility of having *experiences of* the purported existences. Both are certain epistemic thresholds, in one case for a determinate semantic status of claims, in the other for metaphysical respectability.

Now Dummett is without question a great advocate of verificationism, and that his advocacy was forceful in almost all his writings in the philosophy of language. But we can and should, I believe, distinguish the *general form* of Dummett’s criticism of truth-conditional semantic theories as well as his defense for inferential semantics, from the *specific* verificationist slant in the way he carries out the negative criticism.

§2.2.1.a(v) Argument Against Truth-Conditional Semantics

Dummett’s criticism is based on two premises. The first premise, which was never made explicit by Dummett, is that the truth-conditional conception of meaning requires construing the knowledge of what a sentence means in terms of the capacity to *recognize* whether its truth-condition obtains.⁴² The second premise is that certain

⁴² For clear signs that Dummett makes this assumption, see for example the following passage, where, after making the case that *some* of linguistic knowledge has to be implicit, Dummett writes, of the Davidsonian kind of theory of meaning: “The difficulty of giving a suitable explanation of that in which a speaker’s knowledge of the truth-condition of a sentence consists does not lie in deciding what to count as displaying his recognition that that condition is satisfied ... but it is reasonable enough to suppose that, in relation to the speakers of any one language, we can devise a criterion for a speaker’s recognition of the fulfillment of the condition which establishes any given sentence as true.” (*What is a*

Part I: Semantic Preludes

classes of statements are such that there are no procedures that allow one to determine whether the condition of their truth obtains.⁴³ From these two premises Dummett draws the conclusion that, because of the presence of undecidable statements in our language and our apparent ability to understand them, truth-conditional conceptions of meaning fail to give us an account of what our understanding of these statements and hence of our language in general consists in. Clearly then, Dummett's attack on truth-conditional semantics requires an assumption about *how* the manifestation-of-truth-condition-knowledge requirement can be discharged. Dummett's assumption that knowledge of the truth-condition of a sentence can *only* be manifested in one's ability to recognize the truth of that sentence whenever it is true, is blatantly verificationist. Now, Dummett might be right insofar as someone who is committed to explaining meaning in terms of the notion of truth cannot reject the first premise in Dummett's argument, according to which *recognition* of the obtaining of truth-condition is the only way to manifest knowledge of it. But the premise is certainly entirely independent of both the Priority of Understanding Thesis and the Manifestation Thesis. There is no inconsistency, plainly, to agree with Dummett that knowledge of the truth-condition of a sentence has to be manifested, while insisting, *pace* Dummett, that this knowledge results from many small pieces of knowledge – corresponding to words making up that sentence and the way they are strung together – and that each such piece of knowledge is manifested by the kind of practical abilities Dummett himself speaks of. In other words, the specific manifestations of truth-condition knowledge need not be anything on the sentential level, let alone consists in the *recognition* (an epistemic notion) that the given sentence is true.⁴⁴ Dummett's focus on linguistic understanding and insistence that it should not be occult, therefore, is not a manifestation of his verificationist proclivities. It is rather derived from another philosophical orientation.

Theory of Meaning (II), page 45). Here Dummett *takes it for granted* that manifestation of the knowledge of truth-conditions *must* take the form of *recognizing* the truth of the sentence when the truth-condition obtains, or the *criterion* whereby the speaker is capable of such recognition.

⁴³ For discussion on this second premise, especially in relation to Dummett's anti-realism about abilities, see §3.2 below.

⁴⁴ Again, to defuse Dummett's critique the way I am suggesting in the main text might amount to abandoning the idea of basing a theory of meaning on the concept of truth. Still, my aim in the main text is to separate the Manifestation and Priority of Understanding Thesis from the further *verificationist* requirement, rather than to *defend* the truth-based theories of meaning against Dummett's critique.

§2.2.1.a(vi) Late Wittgenstein

That philosophical orientation is that of the late Wittgenstein's, whose *Philosophical Investigations* is relentlessly focused, among other things, on the issue of understanding, including linguistic understanding. What Wittgenstein's answer – to the question about what constitutes understanding an expression – in fact is, is highly disputed, as witnessed not only by ever differing interpretations from able philosophers (Kripke, Wright, McDowell, Anscombe, to name just a few, all have quite substantially, if not radically, different takes on what Wittgenstein recommends), but also by both McDowell's and Dummett's appeal to Wittgenstein in their debate on how linguistic understanding should be explained.⁴⁵ However that may be, it is clear that Wittgenstein rejects any view according to which understanding an expression or an instruction consists in a specific kind of *occurrence*, mental or otherwise.⁴⁶ Wittgenstein's topic and what he says about it in *Philosophical Investigations* make it clear that Dummett, in defending the Priority of Understanding and the Manifestation Thesis, is echoing and pursuing a Wittgensteinian theme, rather than a verificationist one.⁴⁷ In fact, Dummett himself traces his insistence on manifestation of linguistic understanding in practical abilities to Wittgenstein:

The observation that there is no such mental event as a concept's coming to mind is paralleled by Wittgenstein's remark that understanding is not a mental process. One of the advantages of the approach to language as a vehicle of thought [as opposed to the code-conception of language] is that we do not need

⁴⁵ C.f. McDowell's In Defence of Modesty, and Another Plea for Modesty.

⁴⁶ The sections on rule-following in *Philosophical Investigations* are filled with arguments of this kind.

⁴⁷ Dummett's student, Crispin Wright, goes even further and claims that what distinguishes the anti-realist – which is very roughly what Dummett calls "verificationist" – stance is precisely the general demand that all grasp of concepts be capable of being manifested. See his *Wittgenstein On the Foundations of Mathematics*, page 221. But that is, from the point of view developed here, an extremely puzzling claim. For it is one thing to require manifestation, but quite another to stipulate what sorts of things can possibly *count* as manifestation (experience, recognition of the truth of a sentence etc.).

A variant of this point – the separation of the Manifestation Thesis from verificationism or anti-realism – will become extremely important later in my discussion of lawfulness in nature. The point there is that empiricism (which corresponds in the present context to the Manifestation Thesis about linguistic understanding) only links claims of knowledge to experience, but does not dictate what *counts* as experience. The latter is, unlike empiricism itself, not an *a priori* principle, and is only determinable *a posteriori*.

Part I: Semantic Preludes

to look for any *occurrence* save the expression of the thought. ... Wittgenstein said, “To understand the sentence is to understand the language”. He did not mean that (as some American philosophers believe) you would not understand the sentence in the same way if you knew only a fragment of the language to which it belonged. He meant, rather, that, given you understand the *language*, that you are, as it were, in that *state* of understanding, nothing need happen, in which your understanding of the sentence consists, no *act* of understanding, other than your hearing that sentence. This consideration only reinforces our initial idea, that the key to an account of language – and now, it seems, of thought itself – is the explanation of an individual speaker’s mastery of his language.

– Dummett, “*What do I Know When I Know a Language?*”, p.99

To summarize then, Dummett’s motivation for inferential semantics is not a protest against private mental episodes as such, nor a radical empiricist program to do away with the abstract entities *meanings*,⁴⁸ nor a special case of his well-known proclivities for verificationism.⁴⁹ It is rather an advocacy and expansion of a central theme of the late Wittgenstein, according to which the central philosophical question about language is that of understanding, and understanding does not consist in some occult mental occurrence but is “open to (public) view”.

⁴⁸ I am not claiming that all or even the major serious alternatives to broadly use-theoretic semantics endorse meanings as “occult” entities. The kind of truth-conditional semantics initiated by Davidson, and later defended by Evans and McDowell against Dummett, for instance, insists that a theory of meaning does not have referential terms for specific “meanings”, but gives the meanings of expressions indirectly by means of appropriately chosen true basic axioms for a truth-theory in the Tarskian style. In such a theory of meaning, there are no claims of the form “The meaning of the expression E is (=) Θ ” where “E” and “ Θ ” are to be substituted by appropriate singular referential terms of the theory.

⁴⁹ Verificationism and inferential semantics intertwine and interact in much more complex ways in Dummett’s work than I have indicated in the main text. On the one hand, Dummett mobilizes verificationism to remove the major competitor of inferential semantics, that is, to reject truth and truth-conditions as the central concepts for explaining linguistic meaning. On the other hand, Dummett’s argument against classical logic and in favor of intuitionistic logic, which is far more in line with verificationism, is based on the assumption that the meaning of any logical constant, such as the negation sign, is completely determined by its rules of use. That assumption is of course a key tenant of inferential semantics. But as we shall see in our discussion of the concept of harmony (§4.1.4), Dummett’s argument for intuitionistic logic relies *not only* on an inferential conception of meaning for logical constants, but also certain – quasi-verificationist – assumption about the derivability of all use rules from those rules on the verification side.

For Dummett’s verificationist argument against truth and truth-conditions as the central concept of a theory of meaning, see Dummett’s *What is a Theory of Meaning (II)*; for use-theoretic semantic arguments for intuitionistic logic, see his *The Logical Basis of Metaphysics*.

Before turning to the other great early proponent of inferential semantics, Sellars, we need to have a harder look at the Priority of Understanding Thesis, both for understanding Dummett better and for grasping the difference between him and Sellars (as well as Brandom, who is more a follower of Sellars than of Dummett).

§2.2.1.b Priority of Understanding

Despite its centrality to Dummett's semantic thinking, he says surprisingly little in way of justifying the Priority of Understanding Thesis. In this section, I shall (a) consider what Dummett did say about it, (b) consider what Fodor, an arch opponent, has to say *against* it, and (c) sketch an independent defense of the thesis. The reason that an extra defense of the thesis is needed is not because Fodor has good arguments *against* it. In fact, as we shall see, Fodor exhibits profound misunderstandings of Dummett in his criticisms. The reason for an independent defense lies rather in the paucity of what Dummett himself says about it.

§2.2.1.b(i) Dummett in His Own Words

Dummett's most explicit argument for the Priority of Understanding Thesis occurs at the beginning of his *What Is a Theory of Meaning? (I)*:

If the theory of meaning allows us to derive such direct ascriptions of meaning, and if these direct ascriptions are such as to lead in this simple way to a characterization of what it is to know the meaning of each word or sentence in the language, then, indeed, my claim that a theory of meaning must be a theory of understanding is not intended in so strong a sense as to rule out such a theory, merely on the ground that it did not itself employ the notion of knowledge: it would be proper to accept such a theory as being a theory of understanding. If, on the other hand, although the theory of meaning allows the derivation of direct ascriptions of meaning, these ascriptions are so framed as not to permit an immediate characterization of what it is that a person knows when he knows the meaning of a given word or sentence, then, by hypothesis, the theory is inadequate to account for one extremely important type of context in which we are disposed to use the word "meaning".

– *What Is a Theory of Meaning? (I)*, p.3, (underlines mine)

Part I: Semantic Preludes

First note that there seems to be a certain carelessness in Dummett's formulation of the two kinds of cases. He switches from the phrase "what it is to know the meaning of X" in characterizing the first kind of theories to the phrase "what it is that a person knows when he knows the meaning of X" in characterizing the second kind of theories (see the first two underlined phrases). These are not the same ideas, even though Dummett clearly intends to mention the same thing with these two phrases. The first phrase is about the *constitution*, while the second phrase is about the *object* of linguistic knowledge. This is not entirely a lapse on Dummett's part. If it were, it would be surprising because Dummett himself makes this distinction later in the same essay. He criticizes Davidson's truth-conditional theory of meaning on the ground that it merely purports to specify the *object* of linguistic knowledge (namely, that which is expressed by the basic and recursive clauses of a Tarskian theory of truth under appropriate constraints), but not an account of what this knowledge with such an object *consists* in⁵⁰. The reason for the switch of phrases is that, in the dialectic that Dummett develops, he begins by criticizing a theory according to which to know a language is to know truth conditions. But as it turns out, this way of specifying the *object* of linguistic knowledge makes it impossible, in Dummett's view, to answer the more basic question about the *constitution* of such knowledge, so that, according to Dummett, it cannot be right. Besides, Dummett himself does not think addressing the question about the *object* of linguistic knowledge is a helpful way to pursue the constitution question. Clearly, then, what is important for Dummett is that, *if* one undertakes to specify the object of linguistic knowledge in such a way that makes it impossible to explain further the *constitution* of linguistic knowledge, then it cannot be a good theory. The way to read the passage cited, therefore, is to substitute for the second underlined phrase with "characterization of what it is to know the meaning of a given expression".

Having made this slight correction, it is clear that Dummett's argument in this passage, insofar as there is one, is that a theory that does not issue in an account of what it is to know the meaning of any given expression of a language is unable to account for what we are doing when we *attribute linguistic understanding*, or what

⁵⁰ So at the end of the essay, Dummett writes, summarizing his critique against Davidsonian semantics, "...it [i.e. a theory of meaning] must give an explicit account, not only of what anyone must know in order to know the meaning of any given expression, but of what constitutes such knowledge." (Ibid. pg. 22)

sentences of the form “S knows the meaning of X” express. But, so goes this argument, any adequate analysis of a concept must be able to account for all contexts of its use, including the context of knowledge-attribution.

A second argument for the Priority of Understanding Thesis is advanced just a few paragraphs later, where Dummett claims that, for the purpose of having an adequate theory of meaning, it is not only necessary, but also *sufficient*, to have an account of linguistic understanding:

Conversely, it appears to me that once we can say what it is for someone to know a language, in the sense of knowing the meaning of all expressions of the language, then we have essentially solved every problem that can arise concerning meaning. For instance, once we are clear about what it is to know the meaning of an expression, then questions about whether, in such-and-such a case, the meaning of a word has changed can be resolved by asking whether someone who understood the word previously has to acquire new knowledge in order to understand it now.

– *Ibid*, p.4

These two arguments, one for the *necessity*, the other for the *sufficiency* of an account of linguistic understanding for the purpose of constructing a theory of meaning, are not equally important. It is arguably true of *any* kind of entities that, a complete theory about it should enable one to say what it is to *know* one of the kind in question. So the truth of the necessity thesis might appear perfectly general. The sufficiency thesis, however, if it is indeed true, is not clearly true of all kinds of entities. For why should explanation of *one* of many applications of a concept – the application in the epistemic context – suffice to explain *everything* about that concept? Consider for example the color of an object. The explanation of what it is to know that an object has a certain color may be along the lines of: to know the object has a certain color is to be able to use the correct color-word to describe the object after looking at it under normal lighting conditions. But it is far from clear that we have thereby achieved an account of what the color *blue* and the color *red* are, in themselves. We cannot conclude from what we just said about what it is to *know* the color of an object, for example, that an object cannot *be* (purely) blue and (purely) red at the same time.

Part I: Semantic Preludes

§2.2.1.b(ii) Fodor's Ridicule

The sufficiency thesis is at the core of the Priority of Understanding Thesis, and both seem therefore to need some defense. The need is rendered more urgent by the fact that an arch foe of inferential semantic theories, Jerry Fodor, will have no truck with it. Fodor launched a scathing attack in his introduction to the so-called “representational theories of mind”, with a palpable sense of ridicule:

So, to repeat, the methodological doctrine that concept possession is logically prior to concept individuation frequently manifests a preference for an ontology of mental dispositions rather than an ontology of mental particulars. This sort of situation will be familiar to old hands; proposing dispositional analyses in aid of ontological reductions is the method of critical philosophy that Empiricism taught us. If you are down on cats, reduce them to permanent possibilities of sensation ... There is, however, a salient difference between reductionism about cats and reductionism about concepts: perhaps some people think that they *ought* to think that cats are constructs out of possible experiences, but surely nobody actually does think so; one tolerates a little *mauvaise foi* in metaphysics. Apparently, however, lots of people do think that concepts are constructs out of mental (specifically epistemic) capacities. In consequence, and this is a consideration that I take quite seriously, whereas nobody builds biological theories on the assumption that cats are sensations, much of our current cognitive science, and practically all of our current philosophy of mind, is built on the assumption that concepts are capacities.

To sum up so far: it is entirely plausible that a theory of what concepts are must likewise answer the question ‘What is it to have a concept?’ and, *mutatis mutandis*, that a theory of meaning must answer the question ‘What is it to understand a language?’ We’ve been seeing, however, that this untendentious methodological demand often comports with a substantive metaphysical agenda: viz. the reduction of concepts and meanings to epistemic capacities.

– Jerry Fodor, *Concepts: Where Cognitive Science Went Wrong*, pp.3-4, underlines mine.

The ridicule is clearly not meant for theories of concepts only, but theories of meaning and/or contents generally. Indeed, after the quoted passage Fodor continues

to illustrate his general criticism with the specific and “illustrious example” of Michael Dummett and his theory of meaning. I have already expressed skepticism about lumping contents and meaning together in the same category for philosophical investigation. Since contents with mental items as bearers are not the topic of this essay, I will simply focus on Fodor’s critique as one directed at theories of meaning. Fodor makes his complaint plain in his discussion of Dummett, several paragraphs after the passage already quoted:

... the questions with which theories of meaning are primarily concerned are metaphysical rather than epistemic. This is as it should be; understanding what a thing is, is invariably prior to understanding how we know what it is.

– Ibid. p.5

Here we have a flat denial of the Priority of Understanding Thesis. The reason given is simple: the priority of metaphysical explanation over epistemic explanation is general, and there is no reason why it does not hold when it comes to meanings. The passage quoted before this one makes a comparison with cats, presumably to render this point vivid. We come to know what cats are by having sensory experiences about them. But no one, nowadays anyhow, would want to explain cats as abstract constructions out of cat-sensations. Time was indeed such that some philosophers wanted to do just that, but by now everyone has seen, or should have seen, how absurd that idea was.

Is Fodor right to cast aspersions on Dummett’s insistence on the centrality of linguistic understanding? Well, it should not be hard to observe, given what we have seen in our discussion of Dummett so far, that Fodor commits several confusions in his critique. First, Dummett’s emphasis on “practical capacities” that a speaker has to manifest to count as knowing a language does not amount to the privileging of some kind of “*epistemic* capacities”. On the one hand, the practical capacities that Dummett speaks of are not “epistemic”.⁵¹ It is not even clear that they are “mental” capacities. A paradigmatic example of the practical capacities that Dummett has in mind is the capacity to infer “p” from “p and q”. To call that an “epistemic capacity”, is quite a stretch. On the other hand, Dummett has no intention of *reconstructing* meanings out

⁵¹ See more on this point below, where meanings are contrasted with cats and dogs.

Part I: Semantic Preludes

of these or any other practical capacities, the way some early logical positivists wanted to reconstruct cats out of cat-sensations. Manifestation of practical capacities is supposed to be constitutive, not of meanings, but of the *grasp* of what an expression means. Meanings, if there are such things, need neither be constituted by practical capacities, nor by the grasp of meanings. The Priority of Understanding Thesis says what are the most important phenomena to explain, it does not say that everything else ought to be *analyzed by reference to* that phenomena.

The second confusion in Fodor's discussion of Dummett is just that he sees Dummett as driven by *metaphysical* concerns about meanings as abstract entities. This may have to do with the fact that Fodor himself was an erstwhile functionalist, and functionalist approaches are indeed first and foremost driven by metaphysical discomfort. The fact that Dummett happens to be a vocal proponent of verificationism also helps to entrench the impression that, in the question of linguistic meaning, Dummett is similarly motivated by doubts about meanings' ontological status, doubts cast by certain verificationist principles. As I have argued in the last section, however, this is a misreading of Dummett as a meaning theorist. As such, Dummett is as much influenced by Wittgenstein and the latter's single-minded focus on understanding as by his own verificationist tendencies. In fact, nowhere in Dummett's writings on semantics does he suggest *identifying* meanings with sets (or fancier set-theoretic constructions) of practical capacities. He does not say much about the status of meanings as abstract entities at all, let alone link it to the question about linguistic understanding. Fodor's comparison to reconstructing cats from cat-sensations is therefore completely out of place for Dummett, though it might have more bite for other, *functionalist* oriented, inferential theorists of meaning⁵².

§2.2.1.b(iii) More on Fodor's Misunderstanding of Dummett

In concrete philological terms, we can detect Fodor's misleading attribution of metaphysical reduction to Dummett through two discrepancies between Dummett's formulation and Fodor's representation of it. The "epistemic" explanation that Fodor contrasts with "metaphysical" explanation is, according to Fodor, the explanation of

⁵² See for example the discussion below on Sellars.

F) how we know what a certain meaning is.⁵³

Dummett's own formulation of the basic explanatory task for a theory of meaning is to explain

D) what it is to know that an expression means what it means.

The first discrepancy is between the first part of each of the two formulations, and the second discrepancy between the second part of each formulation. First, “how we know...” and “what it is to know...” are quite different sorts of things. For one, it is possible that there are more than one conceivable ways to come to possess a certain piece of knowledge, all of which satisfy the *constitutive standards* specified by answers to the question: what is it to know...? Yet as an empirical fact, *we*, homo sapiens, arrive at the knowledge via *one* of these only. For example, there *might* be intelligent creatures – call them *batmen* for obvious reasons – who come to know the distance and position of an object by sending out ultrasonic waves rather than by *visual* means based on passive reception of photon irradiations. Clearly, the standard for “what it is to know” the distance and position of an object, in order to make out both the batmen and us to be knowers in this regard, must be different from an account of how *we* come to possess such knowledge. Dummett's claim involves, as we have seen, the *constitution* question, not the *empirical* question.⁵⁴ Only by switching from Dummett's constitution formulation to a formulation in terms of empirical epistemic access does Fodor succeed in giving certain semblance of plausibility to his comparison to earlier philosophical projects seeking to reconstruct cats out of “epistemic capacities” for knowing cats.

Let us, therefore, change the first part of Fodor's formulation to the constitutional form “what it is to know...”:

F)* What it is to know what a certain meaning is.

There remains a second discrepancy between Dummett and Fodor's representation of him. The “epistemic” explanation Fodor attributes to Dummett is – after our first

⁵³ I have modeled this on Fodor's formulation quoted above, on page 63.

⁵⁴ Note also that the “how we know X” question could conceivably involve examination of *evidence* for X. But Dummett's way of formulating the main topic of semantics makes evidence the wrong sort of thing to look at.

Part I: Semantic Preludes

correction – the explanation of what it is to know *what a meaning is*. “What a meaning is” can mean either the *identity* of a meaning, or the *kind of thing* meanings generally are. But either way, that is not the object of the knowledge to be explained, according to Dummett. Rather, the knowledge to be explained is the knowledge, for example, *that* ‘La neige es blanc’ means *snow is white*. The latter kind of knowledge is at least in part knowledge about an *expression*, and is *not* explicitly knowledge about an abstract entity called “meaning”. Knowledge of what a meaning is, on the other hand, needs not involve linguistic expressions and is explicitly about a “meaning”. Fodor’s disarmingly simple contrast between “what X is” and “what it is to know what X is”⁵⁵, where “X” can be either replaced by “cats” or “meanings”, therefore, obscures a theoretical decision about whether “means *snow is white*” expresses a *monadic*, or a *relational* property, to be analyzed in terms of a relation to an abstract entity called “meaning”. Only if Dummett had already decided that the latter is the case, would it make sense to attribute to him a desire to do metaphysical reduction about meanings. As a matter of fact, Dummett does not make such a decision, certainly not in his writings urging the centrality of linguistic understanding.⁵⁶

§2.2.1.b(iv) Arguments For The Priority-Of-Understanding Thesis

Fodor’s criticism as well as his diagnosis of Dummett’s insistence of the Priority of Understanding Thesis, therefore, is based on much confusion. But still, is there anything we can say in support of Dummett’s claim of the priority of understanding? I will attempt a sketch of such a defense for the rest of this section.

⁵⁵ Again, after correcting “how we know” to “what it is to know”.

⁵⁶ A comparison that avoids this assumption of relational understanding of meaning statements would be to compare meanings not to cats and dogs, but to some kind of *properties*, say having a specific color. The two explanatory tasks in question would then be (a) what properties of this kind (e.g. having a color) are, and (b) what it is to know *that* an object bears one of these properties. Dummett’s Priority of Understanding Thesis is that, when the kind of properties in question are “means...”, (b) is the more fundamental of the two explanatory tasks. (Note that even if one regards “meaning” as a *gerund* derived from the same verb used to express the property “means *snow is white*”, the phrase “the meaning of ‘snow is white’” still does not express a property. For, though it would not denote an abstract entity, it will denote the *answer* to the question: what does the expression “snow is white” mean?)

For the view that “means *snow is white*” is a *monadic* property, see for example Horwich, *Meaning*. Sellars, whose semantic positions we are about to examine in a minute, also holds the monadic view.

The Most Fundamental Meaning-Related Phenomenon

The first observation from which everything else follows is that the basic distinction between the epistemic and the metaphysical is not adequate for the panoply of phenomena related to linguistic meaning. The distinction is natural for things like cats. There are cats on the one hand, and then there are our beliefs and knowledge about them. It is no doubt also true that they figure in our intentions, we play with them, and they stand in various relations to their non-human environment and things in it. But these can arguably only be properly understood when it has been thoroughly understood what *cats* are. The reason meanings and belief/knowledge about meanings is not a good division of topics is *not* that it leaves out our other, non-epistemic relations to meanings, nor that it leaves out the relations of meanings to other non-semantic entities. Nor is the problem with this division that it takes for granted the existence of “meanings”. Let us grant for the sake of argument that there are entities called “meanings”, *whatever* they are. The problem with the division of topics into meanings and belief/knowledge about meanings is that it misses the most fundamental phenomenon in which meaning plays a role: linguistic practice, linguistic communication in particular. A speech act *is* not a meaning; it *has* a meaning. Nor is it an act of *cognizing meaning*. So speech does not fall into either side of the metaphysical-epistemological division. Nor is speech and communication merely one of many relations meanings can enter into, to us or to other entities. It is in fact the most fundamental phenomena with regard to meaning, because *how* speech is done *completely determines* facts about meaning. It is not as if there were at first meanings, then people figured out a system of speech that, at least approximately, expresses them. If that were the case, it would be hard to explain the variety of languages and the different meanings they are capable of expressing. When we approach a language that we do not know, we are generally prepared to learn expressions that have meanings we are not yet aware of.

One might make the observation that *some* speech is literally “full of sound and fury signifying nothing”, and claim that this shows the relation between meanings and speech is exactly the reverse of what I have claimed: speech acts are *judged* as good or defective according to whether they *have meaning*. The observation is correct as far as it goes, but it does not support the conclusion stated. The proper way to put the observation is that there are constraints on the use of a system of symbols for it to

Part I: Semantic Preludes

be a *language*, in the sense that the symbol types have meaning; there are also constraints on how the meaningful symbol types ought to be used in order for the *tokens* of symbols to be meaningful. But these constraints cannot be thought of in terms of the availability or lack thereof of independently existing meanings for the symbols or their tokens to match onto. If a group of organisms appear to use a set of re-combinable symbols to communicate with each other, it would simply not be intelligible to deny that the symbols have meanings on the separate ground that they do not express *actually* existing meanings. Whether a symbol is meaningful depends only on whether it is used as part of a workable system of communication, *not* on whether it stands in some relation to an element of a fixed set of abstract entities, called “meanings”.

Furthermore, even if it were *technically* possible to collect all the meanings that are expressed by symbols of *any* possible language into a set (let us grant that), it would not make sense to say that this set of meanings *puts constraints* on how speech needs to be like in order to be meaningful. To say that would be like to say that the set of personality profiles obtained by examining every living person puts limits on how the personality of anyone must be if he/she is to be alive.⁵⁷

Argument For The Priority of Understanding Thesis

From this first observation, that linguistic communication is the most fundamental phenomenon for linguistic meaning, it is only a small step to the Priority of Understanding Thesis. “Grasping”, as well as *not* “grasping” the meaning of an expression are, despite the more familiar usage of the verb “grasp” here, *states* rather than acts. Furthermore, to grasp the meanings of expressions in a language is not a state of being *affected* in a certain way. It is rather to possess the *capacity* for activities of some sort. One might say that it is like being fertile or strong rather than being sick or burned. Now it is independently plausible – and this is in fact a version of the *Manifestation Thesis* – that linguistic communication is precisely the kind of *activities* in which the *capacity*, in virtue of the possession of which a speaker counts as “grasping” or understanding meanings, is realized.⁵⁸ If that is so, and if, as we have

⁵⁷ That is, there is merely a *logical* entailment, not an *explanatory relation*.

⁵⁸ In particular, this would guarantee that linguistic understanding lies “open to view” – rather than being something for which *evidence* is needed – *if* we can make sense of directly perceiving that a speaker’s linguistic performances are actualizations of requisite capacities, as opposed to performances

already argued, the activities of linguistic communication constitute the most fundamental phenomena for linguistic meaning, then the following is obviously true. To theorize about what it is to grasp meanings *is* to theorize about what needs to happen for linguistic communication, which consists of the *realizations* of meaning-grasping states, to take place. And the latter, being the most fundamental phenomenon for linguistic meaning, clearly is the object of the central explanatory task for a theory of meaning. Consequently, the question “what it is to grasp meaning” is also the central explanatory topic of semantics.

Appendix to §2.2.1.b: Meaning, Understanding, and Communication

Let me end this discussion by trying to illuminate the unique constellation of meanings, linguistic communication, and the grasp of meanings by contrasting it with the case of ordinary objects like cats and dogs. For it is not enough just to point out how Fodor completely misunderstands Dummett by assimilating his Priority of Understanding Thesis to the phenomenalist strategies for explaining the ontological constitution of cats and dogs. To prevent misunderstandings similar to those of Fodor, it is necessary to make plain the fundamental differences between meanings and cats, so that the futility of an argument based on a supposed parallel between these two cases will be obvious at the very outset.

The first of the two points of contrast I am going to make we have already touched upon in our critical look at Fodor, and it is that, though there is – parallel to cognitive capacities for knowledge about cats – such a thing as cognizing meanings, “grasping meaning” is not it. The second point is that, cognizing meanings, *unlike* cognizing cats and dogs, is a sort of *ancillary* activity, serving the activities of communication.

Let me begin with the first of these two points. The cognitive capacities for cats, like all capacities, are capacities for *actions* or *activities*. In particular, they are *realized* in instances of cognizing cats. But, as we already remarked, the capacity that we attribute as “grasping meaning” is realized in successful instances of communication. Successful instances in communication however are *not* instances of cognizing meanings. This should be abundantly clear to anyone who reflects a bit on how instances of cognizing meanings look like. They are, or result in, propositional attitudes having contents of the form: “somnambulist” means *sleepwalker*; “bank” has more than one meanings; “gavaggai” is meaningless, etc. But surely it is precisely a *mark* of successful communication that, while in

that *happen* to fit the patterns of such actualizations. That is, if we can maintain Dummett’s manifestation requirement while conceding to McDowell that what is manifested is not a *mere* disposition to fit linguistic behavior to patterns, as opposed to activities describable only by giving the *content* of relevant linguistic items, e.g. descriptions like “claiming that p upon seeing ...”, “Infer that p from ...” etc. For the latter point, see McDowell’s *Wittgenstein on Following a Rule*, §11.

Part I: Semantic Preludes

it, we do *not* have such propositional attitudes. In successful communication we notice only facts of the form: S said that p, or simply, p. In short, “grasping meaning” is not a kind of *cognitive capacity for meanings*, and in communication we do *not* stand in an “epistemic relation” to meanings.⁵⁹

The second contrast requires a somewhat lengthier elaboration. The basic point is that cognizing meanings is something that people do *for the sake of* another kind of activity, namely, communication; whereas cognizing dogs and cats needs not serve any extraneous purpose. So, on the one hand, apart from semantic theorizing and certain forms of pastime, explicit meaning cognition does not arise at all in our lives except in cases of breakdowns of ordinary communication, either momentary ones, or semi-permanent ones, as when a speaker is confronted with speech/text in a language that he has not yet mastered. In breakdowns, the first questions one contemplates are not yet directly about meanings, but rather intentional/epistemic ones such as: what does the writer/speaker mean with those words? have I failed to *understand* what she says? Questions about meaningfulness, ambiguity, or word-meaning, then *serve* – together with other considerations – to clarify questions about understanding and semantic intention.⁶⁰ Of the three sorts of things we do – communicating, explicitly questioning understanding and semantic intention, cognizing meaning facts – the last serves the second, and the second serves the first, when it breaks down. In this sense, direct cognitions of meaning facts are ancillary to communication.

With cats, it is widely agreed in epistemology that we do not consider our epistemic relation to them unless our direct cognition of cats breaks down. That appears to parallel what we have said about meanings: we *also* never make understanding or intending of meanings a topic of consideration, unless there are practical difficulties in communication. So if things go well, we notice only facts about cats, dogs, chairs etc. but not our epistemic relation to them. Similarly, the understanding and intending of meaning make themselves noticed only when there is a hitch, so that, when things go well, we only notice that someone said that p. This is true as far as it goes. But first, communication does not *consist* of such cognitions, though it may require them. Second, if there is any thing in communication for cats and dogs to correspond to, they correspond not to meanings of words, but to speech acts that we notice in communication. Though apprehending the words and their meaning is of course crucial if we

⁵⁹ Fodor’s criticism fails on this account alone: to privilege the explanation of what it is to *grasp* meaning is not to privilege the explanation of our *epistemic* relation to meaning.

⁶⁰ “Clarify” either in the sense of *explain*, as well as in the sense of *rectify*. So we might say “‘bank’ has two meanings, and he is using it to mean a financial institute” in the course of figuring out the speaker’s semantic intentions. Again, on the one hand, we might say “She does not know that ‘somnambulist’ means sleepwalker, she thought it means a gymnastic specialist” to *explain* why someone failed to understand our words. On the other hand, we can say *to* the person in question “‘Somnambulist’ means a sleepwalker” to *correct* her understanding.

are to notice in communication facts of the form “S said that p”, such apprehension are not cognitive acts that occur in communication, but, as Dummett says, a *general state* that any competent speaker is in no matter when.

The disanalogy is even more clear if we ask what comes into our cognitive focus when things do *not* go well. In the case of dog/cat observation gone wrong, we would ask questions about epistemic relations explicitly: did I really see a dog? In the case of communications, the questions are: did he understand what you said/wrote? What did she mean with that statement/sentence? It is, as we have been arguing, in order to sort out questions like these that we speak of words meaning this and that, of sentences having the same meaning, of some expression having more than one meanings etc. So direct considerations about meanings *serve* to sort out questions that arise from communication failure, whereas direct cognitions about cats and dogs, *when things go wrong*, is *served by* questions about our epistemic relations to cats and dogs.

To summarize: on one level, the analogy with cats does not work because communication is not direct cognition of meaning-facts, and so does not correspond to direct cognition of cats; on another level, there is an analogy of a sort, but the relation of priority is exactly reversed: direct cognitions of cats are default occurrences, whereas direct cognitions of meaning are triggered by breakdowns of communication

§2.2.2 Sellars: Meaning Nominalism

Sellars is not only probably the earliest proponent of a functionalist approach to semantics, he is also sometimes regarded as the founder of functionalism in the philosophy of mind.⁶¹ Yet I put Sellars under the rubric “functionalist motivation to inferential semantics” not without some hesitation. The first ground for hesitation is that Sellars’ own writings show indubitable influence of Wittgenstein, as one of his seminal papers widely regarded as the beginning of functionalism in semantics is titled “Some Reflections on Language Games”. That essay is basically Sellars’ wrestling with the Wittgensteinian idea of languages as games, a methodological framework made well known by the publication of *Philosophical Investigations* the year before Sellars’ essay. Second, some philosophers see, not unreasonably, the functionalist approach to language already implicit in Wittgenstein’s slogan that meaning is use. Between Sellars and Wittgenstein, there is, in other words, both a historical connection and a connection in substance. Third, and most importantly,

⁶¹ For some historical references, see Dennett’s *The Intentional Stance*, page 341-342, especially the longish footnote.

Part I: Semantic Preludes

Some of what Sellars had to say in semantics – his notion of “conceptual meaning” for example – can be read as directly addressing the question of linguistic *understanding*.

These considerations notwithstanding, there are reasons for identifying Sellars as having, *inter alia*, a functionalist motivation for semantics. Even if one can think of the slogan “meaning is use” as an embryonic form of a functionalist approach to language, it is still the case that there are, objectively speaking, two distinct sets of motivations that can be distinguished: the desire to have an adequate explanation of linguistic understanding vs. the desire to show that there is no ontological problem or mystery about abstract entities called “meanings”. These correspond to two different ways to interpret the slogan “meaning is use”. Dummett takes the slogan to call for a *change of topic* for a theory of meaning from “meanings” to linguistic understanding. But someone more concerned with the status of “meanings” will interpret the slogan as a call to *reduce* “meanings” (meaning is to be reductively analyzed as use) or to *eliminate* them (meaning-claims have no genuine reference to “meanings”). He will not, in other words, read the slogan as recommending a change of topic for theories of meaning. In the case of Sellars, although he has things to say about linguistic understanding, what he specifically emphasizes about the semantic approach he sketches is its *nominalist* character about both meanings as well as symbol-types qua *abstract entities*.

For the sake of clarity, I will postpone the bulk of my discussion of Sellars’ theory about linguistic understanding to §2.4 (specifically §2.4.3 Sellars – Experience & “Conception Of The Norm”) and to §3 (in §3.3.2, on his notion of “conceptual meaning”). Some remarks in that connection, relating to the difficulty to avoid philosophy of mind, will be touched upon below in §2.2.2c The Order of Being vs. The Order of Knowing. Throughout the current section, I shall focus on Sellars’ theory on the status of meanings qua entities.

§2.2.2.a Meaning Nominalism and Functions

The centerpiece of Sellars’ account of the status of meanings qua entities is his analysis⁶² of meaning statements of the form “ ‘und’ (in German) means *and*”. The first thing to note about this analysis is the absolute centrality of it in Sellars’ writings

⁶² See caveat on the word “analysis” below. Here it means merely “a theory about”.

on semantics. It occurs in many different places – in the seminal essay *Meaning as Functional Classification*, in his major work *Science and Metaphysics* (chapter III), in *Language as Thought and as Communication*, to mention a few important ones – and each time with great detail. Moreover, it is hard to explain why Sellars takes it to be the most basic task of semantics to offer an analysis of meaning-statements, unless he views the status of “meaning” as one of the most important questions in semantics. As a matter of fact, Sellars says as much in a reply to Putnam’s commentary on his essay *Meaning as Functional Classification*. By way of pointing out Putnam’s mischaracterization of his approach to a theory of meaning, Sellars writes:

Hilary [Putnam] trivializes the above strategy [for theorizing meaning statements using so-called “functional sortals”] when he suggests that “as long as we have some theory as to what it is for a word to have a particular meaning... we can introduce a sortal ‘·A·’ for all words... with that same normal form”. For it obviously is not a matter of having *some* theory, but of having a *correct* theory, and the strategy I have offered is not something that can be applied *after* a correct theory of meaning has been established, it is (or purports to be) that correct theory itself. For the crucial step in avoiding “meanings as objects” – an objective which Hilary [Putnam] shares – is that of interpreting both the ‘means’ of “‘und’ (in German) means *and*” and the ‘stands for’ of “‘dreieckig’ (in German) stands for triangularity” as specialized forms of the copula. Obviously no progress would have been made if, for example, the functions ascribed to expressions by dot-quoted sortals were those of meaning or standing for specific entities, whether attributes *or* classes, abstract objects or Fregean concepts.

– *Reply to Putnam and Dennett*, pp.458-9, underlines mine.

I shall turn to Sellars’ strategy for analyzing meaning-statements presently. But let me note how (i) Sellars characterizes his analysis of meaning-statements as the “correct theory” of meaning itself, and (ii) he takes one of the goals of his strategy to be nominalism about meanings: to avoid “meanings as objects” by providing a specific analysis for meaning-statements.

I need to mention an important caveat. Though I have been talking about Sellars’ “analysis” of meaning statements, Sellars does not think a statement about functions *analyzes* the meaning of a meaning statement. He takes the relation between

Part I: Semantic Preludes

a statement about meaning and a corresponding statement about functions or roles to be one of “conveying”. I shall examine Sellars’ notion of “convey” in §3.3.3 below, in the context of Sellars’ theory about modal statements. Having made this caveat, I shall continue to speak of Sellars’ “analysis” of meaning-statements, understood simply as his *theory* about them.

To appreciate fully the nominalist character of Sellars’ meaning-statement analysis, we need to look at it more closely. According to Sellars, the three parts of the meaning statement “‘und’ means *and*” are to be thought of this way:

- a) “‘und’” is not a singular term referring to an orthographic type, it is rather what Sellars calls a *distributive singular term*;
- b) “means” does not express a binary relation, it functions rather as a *copula*;
- c) “*and*” does not refer to a “meaning”, it functions as a *functional sortal*.

Each of these claims requires elaboration. A distributive singular term is formed by a sortal with a definite article and it is used not as a definite description, but to make a generic claim. Sellars’ example is “The lion is tawny”. The underlined expression is a distributive singular term because it does not refer to any specific lion. Instead, it combines with a copula and a classificatory term to form a claim that is roughly equivalent with “Lions are tawny”, or “All lions are tawny”.⁶³ What this means is that the original meaning statement makes a generic claim about *tokens* of the type “und”.

The classificatory term in the meaning-statement corresponding to “tawny” is “*and*”, but the latter is not a simple adjective, rather, it is used to classify tokens in German in the following way:

Now it is clearly possible to envisage illustrating sortals which apply to items in any language which (*vis à vis* other expressions in the language to which they belong) function as do the illustrated items in a certain base language...

⁶³ There is obviously the problem that the original generic claim does allow exceptions, such as starved lions or lion-cubs. By contrast, universally quantified claims do not allow exceptions, and it is not clear in what sense a meaning statement allows exceptional tokens. Moreover, the supposed equivalence with universally quantified claims is not enough for what Sellars needs. For, the meaning statement does not apply to *existing* tokens of an expression only. It has a modal force that universally quantified claims simply lack. But I have set these problems aside, as these result from Sellars’ nominalist attitude toward *orthographic types*, something independent of his nominalism about *meanings*. (In addition, the first problem might be avoided if Sellars had chosen something other than generic claims as his model for analyzing meaning-claims; the second problem, that of capturing modal force in a nominalist re-analysis, is quite general for pretty much all such nominalist endeavors).

– Sellars, *Meaning as Functional Classification*, p.427

So the “*and*”-token in the meaning statements “illustrates” tokens of “and” in English, the base language. And “means *and*” classifies certain tokens of German as *functioning the same way as* these illustrated items do, namely, the way tokens of “and” function.

There are two nominalist tendencies in Sellars’ analysis, and only one of them is relevant to our main concern here. Nominalism about *meanings* can be secured by denying that “means” expresses a binary relation (Sellars’ b)&c)). A stricter version of this nominalism will also deny that the monadic predicate “means ...” expresses an abstract property.⁶⁴ But neither nominalization about (monadic) *properties*, nor nominalism about linguistic *types* – which is the drift of Sellars’ a) above – is required to eliminate “meanings” as entities. Meaning nominalism is, in other words, compatible with interpreting “und” as a singular term referring to a linguistic type, *contra* Sellars’ claim a). It helps to distinguish this nominalist position about types from the core of Sellars’ meaning nominalism because it gives a clear sense of the *Zeitgeist* of the era in which Sellars was inventing his semantic functionalism, and the length he was willing to go to meet its nominalistic expectations.

However that may be, the point about Sellars as meaning-nominalist is important because the single most influential (alleged) proponent of Sellars’ ideas in semantics, Brandom, is a *realist* about meanings. This fact can easily obscure the very viable *eliminativist* position about meaning within a broadly inferential framework. In fact, it is the position held by one of the two earliest semantic functionalists. (The other one, Michael Dummett, is most likely a meaning-*agnostic*, as he thinks that all questions about meaning can be answered on the basis of a theory of linguistic understanding.)

§2.2.2.b Functions and Their Pragmatist Interpretation

Now, setting Sellars’ nominalism about types aside, it is still questionable whether Sellars *achieves* his nominalist objectives about *meanings*. We have seen that Sellars aims to “avoid meanings as objects”, and he explicitly rejects *equating* meanings with

⁶⁴ Functional classification (claim (c)) by itself does not give nominalization about the monadic property *means red*, but it very nearly does when combined with a pragmatist interpretation of *functions*. See the next section for more on the latter.

Part I: Semantic Preludes

functions.⁶⁵ For the monadic predicate “means *and*” classifies according to whether the linguistic item to be classified – whether token or type – has a certain *function*, and it is not clear that functions are any more nominalistically permissible than meanings as abstract entities.⁶⁶ So, let us now focus for a moment on functions and roles.

It is difficult to say, from what Sellars has to say in print, why he is more comfortable with functions than with meanings as abstract entities. But it may help to ask why functionalism in the philosophy of mind has the appeal it has. As mentioned at the start of this historical sketch, functionalism promises to reconcile two intuitions: that items in domain F are all there really is, and yet items in a different domain T are both real *and* cannot be reduced to those in F. How does functionalism achieve this? Or at least, how does it give the impression of reconciliation? One answer is that it does that by interpreting the irreducibility of items in T (say mental occurrences) as an irreducible *way of looking at* items in F (e.g. physical occurrences): as playing a certain role in a certain functional system. Insofar as the demarcation and identification of a functional system is presupposed for identification of items in T, the identities of the latter depends on a combination of two factors: how we do the demarcating and identifying of functional systems, and what there is *independent* of our demarcation and identification of functional systems. The latter of these two factors determine that all items in T are also items in F, so that there is nothing that is not an F. The former of these two factors explains why items in T are *not reducible* to items in F.

What I have just sketched I shall call a *pragmatist underpinning* to, or interpretation of, a functionalist theory. I call it a *pragmatist underpinning* to underscore the fact that it explains the functional roles in terms of what *we, as subjects, do* vis-à-vis the functional-role bearers, that is, in terms of our *attitudes* towards them. This is obviously not the only possible take on functionalism and its appeal. But it is pretty much the story Brandom, Sellars’ professed student, tells in *his* version of semantic functionalism; though as usual, Sellars’ story is much more subtle

⁶⁵ See for example, *A Semantical Solution to the Mind-Body Problem* pp. 239-240, and *Some Reflections on Language Games*, page 213. The meaning-statements are said to “convey” information about use-patterns, functions, and so forth, but they do not directly talk *about* these. For more on this, see §3.3.3, especially the discussion on Sellars, below.

⁶⁶ Sellars himself acknowledges this at the end of his paper *Meaning as Functional Classification*, yet he gives no satisfactory responses. The response he *could* have given is sketched in the main text.

and complex. To begin with, although Sellars does give (indirect) signs of approval to this pragmatist interpretation of functionalism, he does not, as far as I know, explicitly tell that story. Sellars' story is also tied up with his complex theory of mind, to which we shall turn shortly.

Let us first look at Sellars' signs of approval for a pragmatist interpretation of functionalism. These signs are to be found most prominently in Sellars' work on epistemology and the philosophy of mind, and in particular, in his well-known attack on the so-called "myth of the given". The myth is, very roughly, the idea that there are items in our mental lives that (i) presuppose no conceptual capacities, and (ii) can play a justificatory role. Sellars thinks that this idea founders on the categorial incongruity between what he calls the "space of reasons" and what he calls the "space of causes". For example, a knowledge claim is in the "space of reasons", while the allegedly epistemic basic sense impressions are in the "space of causes", so that they cannot play the epistemic foundational role after all. But Sellars is decidedly not a *dualist*: he does not think there are two non-intersecting realms of beings, such that items in one can enter into rational or normative relations, while items in the other only enter into causal relations. The way he draws the distinction between "space of reasons" and "space of causes" is well illustrated by the following oft-quoted passage:

The essential point is that in characterizing an episode or a state as that of *knowing*, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says.

– Sellars, *Empiricism and the Philosophy of Mind*, §36
(Underlines are mine, italics Sellars' original)

The "essential point", in other words, is not about what *kind of episodes/states* knowing episodes/states are, independent of our attitudes towards it; it is rather about what it is that we do when we *characterize* something as a knowing episode/state. To *ascribe* the status of being a knowing-episode/state is not a matter of *identifying* the episode/state as being of a certain kind, but a matter of *treating* it in a certain manner, namely, as justifying other knowledge claims. Without delving too deeply into Sellars-exegesis, it seems clear that, in some passages of the EPM, Sellars shows a tendency to treat the status of being in the "space of reasons" as a matter of being *treated by us* as belonging to a network of justifying-relations.

Part I: Semantic Preludes

These are *indirect* indications of Sellars' pragmatist understanding of his functionalist theory of meaning, because the context in which they appear is ostensibly philosophy of mind and epistemology, not theory of meaning. But they have a bearing for us because rules governing the use of linguistic items play a dual role in Sellars' work: in addition to being what has to be grasped for a subject to be credited with *knowledge* of particular matters of fact, these rules are the fabric out of which Sellars' functional system for semantic analysis is constructed.

§2.2.2c *The Order of Being vs. The Order of Knowing*

The issues discussed in *Empiricism and the Philosophy of Mind* are connected with Sellars' theory of meaning in other, more substantive ways. One such connection is relevant primarily for Sellars' philosophy of mind. It is found in Sellars' "psychological nominalism" (the view that "all awareness of abstract entities is a linguistic affair"⁶⁷), and his claim that psychological concepts, such as those of thoughts and sense impressions, are analogically constructed from, therefore secondary to, linguistic concepts.⁶⁸

These (more famous) claims cannot constitute all the connection Sellars thinks there is between language and thought, unless he had maintained that linguistic practice is autonomous and may or may not be accompanied by a dispensable understructure of experiences, beliefs and thoughts. On the contrary, there is evidence that Sellars thought that meaningful *linguistic* discourse is not possible without, among other things, the ability to engage in *mental* activities such as thoughts.⁶⁹

Sellars writes in *Language as Thought and as Communication*, by way of correcting the impression (which might have arisen for his readers) that language and linguistic meaningfulness are phenomena that can exist independent of *mental* capacities, in a kind of summary:

⁶⁷ §29 of *EPM*.

⁶⁸ For a more detail account of the relation between *EPM* and Sellars' *semantic* concerns, especially Sellars' account of linguistic *understanding*, see discussions in §3.2.2.1 below.

⁶⁹ This, in addition to the fact that Sellars does not identify meanings with functional roles, is a second point of departure from Brandom, Sellars' alleged follower in matters semantic. Not only does Brandom identify *specific meanings* with roles in a normative "economy", he also believes to be able to explain *meaningfulness as such* without mentioning anything such as experiences or thoughts. See the discussion to follow for more details. On both points I end up rejecting Brandom's optimistically simplistic revision of Sellars (OSRS for short, and for a revenge of sort of Brandom's fantastic acronyms with which his book is littered).

But above all I should have made it clear that in my view the fundamental concept pertaining to thinking is thinking-out-loud as conceived by our logical behaviorists. This is not to say that I agree with him [i.e. Chisholm] in rejecting the classical conception of thoughts as inner episodes in a non-dispositional sense. Rather I accept mental acts in something like the classical sense, but argue that the concept of such acts is, in a sense I have attempted to clarify, a derivative concept.

– Sellars, *Language as Thought and as Communication*, p.527
(my underline)

Even though the *concept* of mental acts is a derivative concept, the sense of “derivative”, or correlatively the sense of “fundamental”, is, writes Sellars in a footnote to this passage:

[The priority in question, to use Aristotle's distinction, is] in the order of knowing as contrasted with the order of being. As an analogy, notice that concepts pertaining to things as perceived by the senses are prior in the order of knowing to concepts of micro-physical particles, whereas, (for the Scientific Realist) micro-physical particles are prior in the order of being to objects as perceived by the senses

– *Ibid.* (Underline mine)

So in the “order of being” mental acts are still prior to speech, even though, to have the right concepts for *describing* and *cognizing* these mental acts, we must first have the concept of speech acts. This is, according to Sellars, parallel to the fact that, even though electrons and neutrons are, in the “order of being”, prior to tables and billiard balls, to be in a position to *know* about electrons and neutrons, one has to have acquired the concepts of, say, billiard balls, which are required for forming the *concepts* of electrons and neutrons.

These passages are not momentary aberrations. What they express belongs to the most cherished and enduring convictions of Sellars. The views are reiterated in the strongest terms in *Mental Events*, an essay Sellars wrote in the 80s.⁷⁰ What these darkly pregnant passages point to is that Sellars’ semantic functionalism goes well beyond the pragmatist interpretation of norms and functional roles. For, if thoughts

⁷⁰ See §5, 6, and 9 of *Mental Events*.

Part I: Semantic Preludes

are really prior *in the order of being* to speech, then, for the phenomenon of significant speech – beyond mere sound and fury – to emerge, it is not enough (as a *pure* pragmatist functionalist might have held) for there to be sound bites that are, but might not have been, *treated by us* as part of a functional system of sounds, whose bits are, but might not have been, *treated by us* as being related to each other normatively. For, “in the order of being”, there have to be, *in some sense, thoughts* before there is (or, *if* there is to be) meaningful speech. Thoughts do not, though the *concept* of thoughts does, depend on our concept of speech. We conceptualize about thought with the help of the analogy to the concept of speech, which has a functionalist dimension that can be given a pragmatist underpinning.

The pragmatist underpinning of functionalism is therefore only one strand in Sellars thinking about linguistic meaning. Another strand is the insistence that thoughts are prior to speech *in the order of being*. I will not stop to consider how these two strands fit each other and with the rest of Sellars’ semantic conception, as that is an ambitious topic going beyond the modest aims of this portion of my essay. We will see presently, however, that this pragmatist-functionalism will re-appear in a single-minded form in Robert Brandom’s work on semantics.⁷¹

§2.3 Content of Thoughts vs. Meaning of Expressions

§2.3.1 Varieties of Positions

Before getting into more details about different kinds of inferential semantic theories, in particular that of Brandom’s, let me turn briefly to an issue that I promised to deal with at the beginning of this chapter. There I said that “content” and “meaning” are different categories, so much so that there may not be a generic concept encompassing both. One way to see this difference is to note that Dummett’s way of posing the central question of semantics – what it is to grasp the meaning of an expression – cannot be meaningfully posed for “contents”. As a matter of fact, this issue is a key fault-line dividing the functionalist motivation to inferential semantics on the one hand, and the Wittgensteinian-Dummettian quest for linguistic understanding on the other: whereas Dummett and the late Wittgenstein’s focus is on

⁷¹ A similarly one-sided influence can be found in the writings on the philosophy of mind by another student of Sellars Daniel Dennett.

languages, the functionalists have often considered *mental states or acts* as the primary content-bearers.

The functionalist approach to content/meaning is historically closely aligned with the program to make sense of the *ontological status* of, not just linguistic, but above all *mental* items, on the premise that the world describable by physics is all there is. For most varieties of inferential semantics motivated this way are explicitly and primarily theories about the representational content of mental states. In fact, some of these writers also identify theories about the content of mental states as *part* of the functionalism about mental items.⁷²

For the rest of this essay, I shall limit my attention to meanings of linguistic expressions. To do that is not only terminological prudence, or justified on the ground of ordinary linguistic practice alone. Nor is it entirely grounded in the pragmatic necessity to exclude topics in the philosophy of mind. One important factor here is that I do not want to foreclose the possibility that there are in fact two *substantially different* concepts in the neighborhood. In fact, I think arguments can be made that “meaning” is an externalist concept while “representational content” has an internalist component, capturing what Fregeans call “cognitive significance”. The present essay is however not the place to carry out such arguments. Some of considerations in favor a conceptual distinctions are found in the appendix below.

Appendix to §2.3.1

Now, apart from historical, and so on the face it contingent alliances, are there arguments of substance in support of keeping the concepts of meaning and content separate? One conceivable argument is as follows. Though I will make it by focusing on Dummett’s insistence on the *grasp of meaning* as a central topic of the philosophy of meaning, Dummett himself does not make this argument – because he was explicit *ab initio* that the object of his investigation is language. The argument is that the very talk of *grasping meaning* makes sense only if the meaningful items in question are expressions of a language. There are two reasons for this. First of all, it really does not make sense to speak of the “meaning” – in the semantic sense of the word – of a thought (a thinking episode), or a belief (a believing state). Yet some have employed the expression in just this way.⁷³ The appropriate vocabulary for thoughts and beliefs would be “content”, for example, in the sense of representational

⁷² For a representative example of this, see the passage from Ned Block below, in §2.4.1.1.

⁷³ C.f. for example Sellars’ *A Semantical Solution to the Mind-Body Problem*, where he speaks of the “meaning” of a state that can be identified as mental.

Part I: Semantic Preludes

content⁷⁴. Contemporary literature tends to use “content” and “meaning” interchangeably, or to use “content” as if it were a generic concept encompassing both linguistic meaning as well as representational content of a mental state, or, in more rare cases, to use “meaning” as a generic concept⁷⁵. Yet, judging from the way we ordinarily use both expressions, there is a *prima facie* reason to keep the two apart.

The second reason that “grasping meaning” does not apply to thoughts and other mental items is that it makes little sense to speak of “*grasping*” the content of a mental episode or state. To see this, consider how the following question might be answered: what is it for me to grasp the content of a thought? There are two cases: either the thought in question is mine, or it is someone else’s. If the thought is mine, then the question will seem puzzling: there is nothing more to *grasping* the meaning or content of one’s own thoughts than simply *having* these thoughts, or having them *occur* to me. For it is not intelligible to have a thought but not understand its meaning or content. The oft-repeated point of semantic externalism, that in stating what I believe, I may not know the meaning of some of the words I use, has no application here, *unless*, I suppose, attribution of thoughts is literally the attribution of silent inner-speech. For the point is not that I must be able to *state* knowledgeably what I believe or think, nor that I must know what I *say*. The point is rather that, if somehow it is correct to say that I do not completely grasp a thought attributed to me, then the *attribution* of that thought to me is problematic to begin with. The fault, one might say, does not lie in me, the putative thinker, but lies in the attributor. On the other hand, if the thought in question is not mine, then the question “what it is to grasp the content of that thought?” cannot be answered unless it first be answered how I *perceive*, or have epistemic access in some other way, that *foreign* thought. If, for example, I perceive someone else’s thought by hearing his words, then the answer to the question can simply be that I grasp the content of his thought by grasping the meaning of the words he utters. In either case, the point boils down to the following two points. First, quite generally, it does not make sense to say “grasping the meaning/content of X” unless we can distinguish between *merely* perceiving X but not perceiving its content/meaning on the one hand and *not* perceiving X at all, on the other hand. In other words, to grasp the content of X is to discern a specific *abstract character* of X, which is something that one must be able to *perceive independently*. Second, unlike in the case of a linguistic token, which one can perceive without knowing its meaning, there is simply no intermediate epistemic status of perceiving a thought, but not its content, and this is the case

⁷⁴ That is, what the subject represents with, or through, the given thought or belief, to herself, about the world. But there are other notions of contents: information content, cognitive content, or propositional content of any attitude that can be characterized with a “that”-clause.

⁷⁵ See the earlier footnote on Sellars.

both when the thought is mine *and* when the thought is foreign. Sure one can perceive *that* S is having *some* thought, but that is not quite the same thing as perceiving the thought S is having.⁷⁶

*Meaning, Content, and Cognitive Significance*⁷⁷

The difference of the two categories in question is ultimately due to the fact that linguistic expressions serve *communication*, and are *intended to be understood by others* (or the later self), whereas mental representations serve, if that is the right way of putting it at all, “oneself”. This fundamental, and one might say *functional*, difference between linguistic expressions and representational states/acts means that any difference in cognitive significance is an unintended imperfection of language, for it hampers rather than enhances communication. Theoretical question about meanings, as I have argued earlier, arises originally from frustrated communication; so to *answer* a meaning question is part of an effort to clear any obstacles to communication, and therefore the satisfactory answer has a natural tendency to be neutral about cognitive significance. By contrast, the attribution of content to cognitive states and thoughts occurs not to resolve problems of communication, but in the context of *psychological explanation*, for which cognitive significance is of the utmost importance.⁷⁸

⁷⁶ Davidson has made somewhat similar though not identical arguments for a different point: that propositions, as contents of thoughts and other “propositional attitudes”, are not the psychological or epistemic *objects* of these attitudes. But like practically everyone else, he treats “propositions” as candidates for characterizing (“specifies”) an attitude *and* for representing what a sentence means (though he ultimately rejects them, of course). Certainly he does not argue for a distinction between the content of attitudes and the meaning of sentences. For details, see Donald Davidson’s *What is Present to the Mind?* 1991.

⁷⁷ For the view that cognitive significance need not to belong to meaning, see also discussions in §1, especially on Kripke’s lectures on proper names in the appendix to §1.3.2.

⁷⁸ The above line of consideration based on the “functions” of expressions and mental representations obviously assumes the centrality of communication as a function of language. Against that, it can be maintained, for example, that differences of cognitive significance are useful, if a speaker’s goal is to mislead or to manipulate, rather than to communicate. But manipulations are not what language is *for*, we might say, and linguistic manipulations are not so much use as misuse of language. Even apes manipulate and mislead each other, and it is done by humans with non-linguistic means all the time. Language can be used for other, and better, things as well, such as entertainment, from the cabaret to the poetry of Rilke and prose of Joyce, and that function is at least partly based on differences in cognitive significance, among expressions that have “equivalent” powers of communication. But again literature and literary entertainment only makes use of differences in experiential perspectives crystallized in words the way movies and paintings make use of these differences in another medium. Language did not emerge to meet the need for a special sort of entertainment. Rather, artistic impulse is a universal force, and tends to leave no exposed area uncovered in the vast terrain of experiences, though no area is exposed *so that* it may be put to artistic use.

§2.3.2 Taking Stock

We have seen how functionalist motivated semantic functionalism differs from Dummett's version motivated by quests after linguistic understanding. For *us*, this gives rise to two of the four demarcation parameters mentioned earlier ((C1) and (C2)): whether we go with Dummett and set the explanatory goal to be linguistic *understanding* as opposed to other meaning-facts, and whether we ought to employ the functionalist strategy, as Sellars does, to give an account of meanings as entities. A third parameter ((C3) above) has already been discussed and explained: I shall concentrate on the *linguistic* concept of meaning rather than dealing with a generic concept of content.

The issue about linguistic understanding (C1) is tied up with the fourth, last choice we mentioned, about the right kind of semantic explanation ((C4)), which is the topic of the next section. So I will postpone discussions on it. But the issue of the ontological status of "meanings" ((C2)) can be dealt with straight away: I will not further pursue the Sellarsian project of trying to eliminate apparent reference to meanings in our ordinary meaning-talk by a re-analysis of such talk combined with a functionalist construal of linguistic activities. The reason, apart from the limited scope of this essay, is that I am simply not compelled by the metaphysical worry about the status of *meanings* as entities. The reason is not just anti-eliminativism about meanings. To explain, it is helpful to distinguish two issues around the topic of meanings as entities. The first is the question of factual linguistic usage of the word "meaning": is it plausible to construe at least some of our uses of "meaning" as referential? The second is the question of the theoretical utility of abstract entities in a semantic theory, whatever they are called in our meta-theory, "meanings" or "functions". Sellars may or may not be right about the *first* question. But even if he is right, in that our meaning-talk is best *not* construed as referential, there remains the possibility that *some* kind of abstract entities can play constructive roles in semantic theories. It is the second of these two issues – rather than the semantic analysis of meaning-talk – that will interest us in this essay. More specifically, what is of interest is the *explanatory* role played by functions or roles of linguistic expressions. What is the best way to analyze meaning statements just does not interest us here.

Now I have argued earlier in this essay that some abstract devices may have *non-explanatory* roles in semantic theory, in that they allow us to *model* such

semantic phenomena as communication and entailment relations. Possible-worlds for example are valuable precisely for such a modeling role. For such modeling purposes, it is mere idle speculation to ask what is the nature of possible-worlds. For *explanatory* purpose, on the other hand, it is not an idle question to ask what functions are. This is the case especially the explanation aimed at is what I shall call *ti-esti* explanation of linguistic meaningfulness. To this I now turn.

§2.4 Which Semantic Functionalism?

§2.4.1 Two Kinds of Semantic Explanations.

We have seen in §2.2 that late-Wittgenstein-Dummett style semantics declares a change of topic from “meaning” to *grasping* meaning, while Sellars-Functionalism style semantics was motivated by the promise of the functionalist strategy to assuage certain worries about the ontological status of *meanings as entities*. But functionalism, as a *form* of explanation, characterizes both styles of semantics. The difference is just that Dummett aims to explain, not straightforward facts of meaning, but facts of linguistic understanding, and to do that in terms of the role of an expression in *one individual’s* linguistic activities.⁷⁹

So which sort of semantic functionalism is the right one? Or, less ambitiously, which sort of semantic functionalism do we want to adopt in the context of this essay? Comparisons of different varieties of semantic functionalism can be facilitated by a distinction between two different sorts of semantic explanations. And to this I now turn.

I will first explain the distinction against the background of a theory about strictly meaning-facts, rather than theories directly about linguistic understanding. For this kind of theory, we can distinguish at least two different things the theory can aim to explain. The first kind of explanation aims at what might be called the *ti-esti*, or the *constitutional* question of meaningfulness: what does it consist in for something to have meaning at all, in the way a linguistic expression has meaning, regardless *which* meaning it has? The second kind of explanation aims at what might be called the *identity question* about meaning: in virtue of what do various expressions have the

⁷⁹ For more on the sense in which Dummett can be thought of as a functionalist, see the section on Dummett below.

Part I: Semantic Preludes

various meanings they have? The distinction is neutral to whether one is a nominalist or realist about meanings, that is, neutral to how one analyzes “has the meaning...” or “has meaning”.

The distinction can be illustrated by analogy with another contrast that is perhaps more evident. An interesting and difficult explanatory task aims to clarify what it is to be human, the *concept* of a human being, if you will. Any acceptable answer would have to, it seems, mention rationality in some form or other. But rationality needs not play a role at all in the explanation of what gives you and me *distinctive identities*. One such explanation would point to the material origins of us: the different eggs and sperms from which we came to be. The *ti-esti* and the identity questions about being human, in other words, may be best answered by mentioning completely different sorts of things.

In fact, the distinction I just drew has been, at least implicitly, acknowledged by philosophers of language. A prominent example of that is Ned Block’s overview essay on the so-called “conceptual role semantics”:

One major motivation for CRS [conceptual role semantics] is a functionalist (q.v.) approach to the mind generally. Functionalism says that what makes a state a mental state is the role it plays in interacting with other mental states in a creature's psychology. This gives rise to a weak form of CRS: a state is *meaningful* (i.e. has *some meaning or other*) by virtue of the fact that it plays a certain role in a person's psychology. ... The weak form of CRS is functionalist about mental *states* (e.g. beliefs vs. desires) without commitment to functionalism about mental *contents* (believing *that snow melts* vs. believing *that grass grows*). The former claims, in effect, that *having any content at all* depends upon having a role in certain processes. The latter, along with CRS, claims that the *specific content* a state has depends upon its role. The functional role of a thought includes all sorts of causes and effects that are non-semantic, e.g. perhaps happy thoughts can bolster one's immunity, promoting good health. Conceptual roles are functional roles minus such non semantic causes and effects.

– Ned Block, *The Routledge Encyclopedia of Philosophy*

Bracketing my reservation about applying terms like “semantics” or “meaning” to mental items, the passage cited distinguishes two ways of applying the functionalist strategy. I shall call one of them *generic functionalist explanation* and the other *specific functionalist explanation*. A *generic* explanation seeks to say what it is to *be* an item of a certain kind, and a *generic functionalist* explanation says that to be an item of a certain kind is to play *some* role in a functionalist system. A *specific* explanation seeks to say what it is for items of the given kind to have a specific quality-kind that characterizes all items of that sort (so having meaning is the *kind* of qualities that characterizes all linguistic expressions, and having a particular meaning is a specific quality of this kind). A *specific* functionalist explanation explains the possession of a specific quality in terms of the items playing a *specific* functional role in the functionalist system used for the generic explanation.

Formulated this way, a specific functionalist explanation is still a *constitution* explanation, and the specificity of its explanandum makes it, as Block puts it, stronger than the generic functionalist explanation. Specific functionalist explanation is, in other words, a specific constitutional explanation. We can think of such an explanation as having two parts: a generic constitutional explanation (playing *some* role in a certain functional system) plus an *identity* explanation (the role it plays is *this* role).⁸⁰ Now generalizing this thought, we can make our earlier distinction more precise by distinguishing the *identity constitution question* from a non-constitutional *identity question*, with the latter being the differential between a *ti-esti* question and a corresponding identity constitution question. More formally:

(Identity Constitutional Question):

In virtue of what does *E* mean what it means?

=

(*Ti-Esti* Question): In virtue of what does *E* have a (linguistic) meaning at all?

+

(Identity Question): What makes it that *E* means what it means?

⁸⁰ Another place where something *like* my distinction can be found is right in Brandom’s *Making It Explicit*, in the form of “strong” inferentialism vs. “weak” inferentialism (page 131-2). But Brandom’s distinction is not quite the same as mine, for it has to do with whether inferential articulation is sufficient or necessary for “conceptual contentfulness”. Since it is left open whether in this formulation “contentfulness” is to be read as contentfulness *as such*, or as the possession of a *particular content*, the distinction intended cannot be about this unspecified contrast.

Part I: Semantic Preludes

The equation means that the identity constitutional question is to be answered by answering a generic constitutional question and then answering an identity question, which is no longer constitutionally formulated. The contrast we want to draw, then, is between the *ti-esti* question and the (non-constitutional) identity question.

What makes the *ti-esti* question radically different from the identity question is that the former is far more difficult to answer than the latter. Later in the chapter I shall discuss what an answer to the identity question can look like. In the rest of this section I will illustrate the difficulty of the *ti-esti* question with the examples of Dummett, Sellars, and Brandom.

The idea expressed by Block's passage just quoted, that to be contentful is to play *some* role in a functional system is the general form of functional semanticists' answers to the *ti-esti* question. What further characterizes Dummett, Sellars, and Brandom is their appeal to a functional system of *linguistic use* – rather than, say, causal system of neurological states. The success of a use-functionalist account of the *ti-esti* question depends on at least two factors. The first factor has to do with the criterion of adequacy that, any account of meaningfulness must make out the use of a linguistic expression – say the utterance of a sentence – as *done with understanding*. I shall call this the Understanding Criterion:

(Understanding Criterion) In order to be successful, an account of the *ti-esti* of linguistic meaningfulness must make out linguistic actions as actions done with *understanding*.

The second factor has to do with the *representational dimension* of linguistic actions: unlike other sorts of actions done with understanding, linguistic actions *purport to represent how the world is*. The second criterion of adequacy is therefore⁸¹:

(Representation Criterion) In order to be successful, an account of the *ti-esti* of linguistic meaningfulness must make out linguistic actions as being *about the world*.

⁸¹ Ede Zimmermann suggests that these two criteria for *actions* to possess linguistic meaning correspond to two aspects of *linguistic tokens*: sense and reference (as conceived by Frege). To properly produce a linguistic token having reference, one has to be *representing* something in the world, and more generally, to properly produce any linguistic token, one has to *understand* the significance (“sense”) of that token.

Both criteria turn out to be extremely difficult to meet. For, to answer the *ti-esti* question of meaning, a use-functionalist explanation is under pressure to make do with *thin* descriptions of putative functional role bearers. They are, for example, under pressure to start with such items of “mere vocalizations”, rather than assertions, the concept of which already implicates understanding and representation. By adopting rich descriptions the explanation would threaten to be empty. But adopting thin descriptions as the starting point can frequently make functionalist boot-strapping seem an extremely daunting undertaking. One possibility would be to start the account with a notion of pre-linguistic intentionality and mental capacities, and build out of these and linguistic practice a full-blown notion of linguistic understanding and representation. None of the three theorists discussed here – Dummett, Sellars, and Brandom – can warm up to such an idea of a language-independent intentionality however. In any event, to answer the *ti-esti* question, it appears that decisions have to be made about a nexus of issues concerning intentionality in general, and about the role of the mental in particular. And if the view is that the mental does play a role in the explanation of what it is for something to be meaningful, then a story has to be told about that role.

It is instructive to note that, by contrast, there are no such Charybdis (empty explanation) and Scylla (impossible bootstrapping) to steer clear of for the *identity* question of meaning. Again, possible forms of identity explanations are discussed later in the chapter.

In the following, I shall illustrate how Dummett and Sellars had difficulty meeting the Understanding Criterion, and Brandom with both, but especially how he fails the Representation Criterion.

§2.4.2 Dummett – Implicit Knowledge/Consciousness

Dummett, as we recall, wants to explain linguistic understanding in terms of “practical abilities” of sorts. So the *ti-esti* explanation for him is an account of *what it is to understand a linguistic expression (in any given way)*. Dummett’s strategy of analyzing understanding in terms of capacities to do things with words, specifically to make inferences with them, can be thought of as a functional explanation as well. In this explanation, the functional system consists of various linguistic actions the putative speaker is capable of performing: making inferences of a certain form to and

Part I: Semantic Preludes

from sentences of a certain form. For a vocalization to be done with understanding, then, is for it to stand in such inferential relations to other phonetically related vocalizations by the same putative speaker.

Dummett recognizes however that linguistic understanding may not be analyzed without residue in terms of mere capacities for vocalizations, however complex. He points out, for example, that there is a difference between the abilities involved in riding a bicycle or swimming on the one hand and the abilities involved in baking a cake by a recipe, where the ability to read and understand the recipe is required. Dummett identifies this dimension of variation as “degrees of consciousness with which a person may perform a skilled operation.”⁸² He makes the incisive observation that there is no such thing as seeming or pretending to swim: if someone seems to be swimming he *is* swimming. But it is perfectly intelligible for someone/something to *seem* to speak a language without actually knowing how to speak it.⁸³ The skilled activities required for speaking a language are, in other words, not purely practical like that for swimming; they need an accompanying “consciousness” of what one is doing.⁸⁴ At this point, however, Dummett’s story begins to waver. It is not clear what this accompanying consciousness amounts to. At the end of the essay in which he deals with the nature of the practical abilities involved in speaking a language, Dummett reverts to the term “knowledge”:

He [the speaker] can make use [that is, consciously] only of those regularities [of speaking] of which he may be said to be in some degree aware; those, namely, of which he has at least implicit knowledge.

– Dummett, *What Do I Know When I Know A Language?* p.105

But the very explanatory task formulated at the beginning of Dummett’s essay was to say what distinguishes the kind of “practical knowledge” involved in speaking a language from the kind of “practical knowledge” involved in swimming. By

⁸² Dummett, *What do I Know When I Know a Language?* page 95.

⁸³ *Ibid.*

⁸⁴ This argument would be challenged by someone who believes in the Touring Test. I will let this objection pass. The debates would be complex. I note simply (a) both Dummett and Sellars recognizes the need for some sort of reflective awareness for linguistic understanding, and (b) McDowell as well as Sellars argue, in their different ways, that world-directedness required for intelligent language use cannot be achieved without the mediation of *experiences*. Both these issues will be discussed.

characterizing the difference in terms of “implicit knowledge”, no progress has been made. For, if the difference lies in the *implicitness* of knowledge, then, because any practical knowledge is implicit, one has to say what distinguishes the implicitness of *linguistic* knowledge from the implicitness of the knowledge for *swimming*. If the explanation is rather that linguistic knowledge is a more robust, or genuine kind of *knowledge*, perhaps because of the presence of *understanding*, then we have merely *restated* what we set out to explain – that unlike other skilled performances, linguistic performances are done with *understanding* – rather than explaining it.

It remains open for us to say that linguistic knowledge is *sui generis*, and one can at best *identify* it, but one cannot explain it in terms of other varieties of knowledge or capacities. In fact, something like that is precisely what Dummett’s arch-foe in matters of semantics, John McDowell, holds.⁸⁵ The problem for Dummett to take on such a view is that he would be giving up the explanatory undertaking of linguistic understanding via the *Manifestation Thesis*, according to which linguistic understanding is constituted by “practical abilities” to do things describable in terms *not* involving the concept of understanding.

§2.4.3 Sellars – Experience & “Conception Of The Norm”

As we saw, Sellars’ main interest in meanings takes the form of wanting to eliminate them. What remains is a functional description of expressions in terms of their “roles” in a language game. After “meanings” are gone, there remains the question: can Sellars make out items in this description of the linguistic practice really *speech*? in particular, as occurring with *understanding*? For Sellars, the “roles” of an expression

⁸⁵ C.f. for example, McDowell’s *In Defense of Modesty*, page 99:

“Rejecting psychologism is taking the view that the senses of utterances are not hidden behind them, but lie open to view: that is, that to be a speaker of language is to be capable of putting one’s thoughts into one’s words, where others can hear and see them. One great beauty of those modest theories of meaning that are “homophonic” is the distance they go towards making that idea unproblematic, by showing that we need not think of it as amounting to more than this: the thought (say) that some table-tops are square can be heard or seen in the words “Some table-tops are square”, by people who would be able to put their own minds into those words if they had occasion to do so.” (*Underlines mine.*)

This is the strategy of letting mind permeate the performances that are supposed to manifest a certain capacity of the mind. And it comes up again in McDowell’s interpretation of Wittgenstein:

Wittgenstein’s problem was to explain how understanding can be other than interpretation (see §7 above). This non-‘anti-realist’ conception of a linguistic community gives us a genuine right to the following answer: shared command of a language equips us to know one another’s meaning without needing to arrive at that knowledge by interpretation, because it equips us to hear someone else’s meaning in his words.

Part I: Semantic Preludes

are characterized in terms of use *in accordance with inference rules*. This is the feature that underlines Sellars' account of linguistic understanding. For as we shall see later in §3, for Sellars, to grasp a concept is to learn to use the corresponding expression in accordance with certain (e.g. material) inference rules. So it is broadly similar to Dummett's story. Whether Sellars has more success, depends on the details on his account.

At first sight, it may appear that function-bearers in Sellars' language game are describable in rather thin terms (perhaps along the lines of "making an utterance ...") without the vocabulary of understanding. Here is a characteristic passage in *Meaning as Functional Classification*, where he explicitly argues for replacing the talk about meanings with functions of words:

The difference [between thinking-out-loud, which requires "knowing the meaning of words", and parroting words] is rather that the utterances one makes [in thinking-out-loud] cohere with each other and with the context in which they occur in a way which is absent in mere parroting. Furthermore, the relevant sense₁ of 'knowing the meaning of words' (which is a form of what Ryle has called *knowing how*), must be carefully distinguished from knowing the meaning of words in the sense₂ of being able to talk about them as a lexicographer might – thus, defining them. Mastery of the language involves the latter as well as the former ability. Indeed they are both forms of know how, but at different levels – one at the 'object language level, the other at the 'meta-language' level.

– pp. 430-31 (*Underlines and subscripts mine*)

The characterization "cohere with each other and with context" is really just an empty place-holder, and does not add any substantial requirement beyond the exhibition of regularities. The second part of this passage is slightly more illuminating: in contrast to Dummett, Sellars here seems to think that the requirement of understanding does *not* entail the need for any additional description of the mastery of language beyond that of "know how". Unless with "cohere" Sellars means to allude to some story that goes beyond conformance to patterns and rules (of correctness), he appears to have a functional description about language with a *thin* basis (i.e. mere productions of

sound-sequences), and according to which “accompanied by/with understanding” is just a matter of these sound-productions fitting some pattern.

But Sellars does talk of “ought-to-do” and “ought-to-be” norms, applicable to language *trainers* and language *learners* respectively. In other words, Sellars’ description of a functional system is not purely naturalistic but also employs *normative* vocabulary. The question now appears to be: does Sellars’ use of *normative* vocabulary suffice to make out our language game to be a *rational* practice, a practice whose elements are done with *understanding*? Probably not. The problem is that, given Sellars’ *pragmatist* understanding of normative statuses, the talk of norms does not fundamentally alter the physicalist nature of the functional system. In particular, “ought-to-be” norms appear to an evaluator who *takes* certain linguistic performances *as* subject to correctness norms. “Ought-to-do” norms appear to an evaluator *of an evaluator* because, presumably, he *takes* the latter’s performance to be subject to rules of action, properly so-called.

Moreover, the distinction of two kinds of norms does not take back this pragmatist interpretation. Sellars insists that the difference between the two kinds of norms is only one of perspective: when I am *evaluating and correcting* another speaker, my activities are subject to ought-to-do norms, but when my speech is *being evaluated and corrected*, my activities are subject to ought-to-be norms. This insistence is simply part and parcel with Sellars’ denial that there is a difference in kind between the two sorts of activities he distinguished. Both are performances of the same “know how”, which is just the ability to produce sound-sequences that conform to *some* correctness norms (albeit very complicated ones). There is, it seems, no place for understanding in Sellars’ functional system after all.

So far, I have mentioned elements in Sellars’ story that also appear in Brandom’s: norms, and a pragmatist interpretation of these. But normativity and pragmatist underpinning of functionalism is only one set of characteristics of Sellars’ semantic functionalism. As we already noted⁸⁶, Sellars’ theory is complex also because of the way it meshes with his philosophy of mind, especially because his insistence on the priority – in the “order of being” – of thought and other mental items over speech.

⁸⁶ In §2.2.2c The Order of Being vs. The Order of Knowing above.

Part I: Semantic Preludes

It is indeed a great difficulty attending Sellars interpretation to figure out *precisely how* these two aspects mesh with each other. That the mental *does* have a role in Sellars' theory of meaning, especially in making out our linguistic moves to be *rational*, is quite clear. This not only manifest itself in Sellars Kant interpretation in *Science and Metaphysics* and in his "epistemological" discussions⁸⁷ in *EPM*. Sellars developed an elaborate theory of *experience* in these places. Importance of the mental – this time concerning the intra-linguistic *norms* – for linguistic understanding is also found in a passage in *Inference and Meaning*, a relatively early essay published two decades before *Meaning as Functional Classification*. In that paper, Sellars clearly recognizes that the elements of his semantic functional system, namely instances of *rule-governed behavior*, cannot be simply characterized as naturalistic elements fitting certain external norms:

A uniformity in behaviour is rule-governed not *qua* uniformity, for then all habitual responses would be obeyings of rules – which is clearly not the case – but *qua* occurring, in a sense by no means easy to define, because of the conception of the norm enjoined by the rule. Yet the fact that both rule-governed and merely associative uniformities are *learned* uniformities, and differ in this respect from, say, the uniformities studied in chemistry, has blinded many philosophers to the important respects in which they differ from one another...

– Sellars, *Inference and Meaning*, p.284

So rule-governed behavior, which is what utterances of sound-sequences must be if they are to be significant speech, is more than behavior *conforming* to some pattern or norm. Where Dummett has "consciousness" of patterns and "reflections" – for distinguishing significant speech from simpler kinds of skillful performances – Sellars has "conception of the norm". Dummett falls back to "implicit knowledge" while trying to explain what the differentiating character of consciousness or reflection after all is, whereas Sellars admits straight away that the "conception of the norm" he alludes to is "by no means easy to define".

⁸⁷ Which are, as McDowell rightly points out (for example in his Woodbridge Lectures), really about intentionality, not about empirical *knowledge* per se.

All in all, Sellars' "physicalist" functionalism – where the potential functional role bearers are specified physicalistically and the normative relations eventually are explained in pragmatist terms through *attitudes* – is complemented both by a sophisticated theory of *experience*, which plays a role in making out language *entry* moves to be rational, and a recognition, similar to that of Dummett's, that *intra-*linguistic moves are accompanied by a certain "conception of the norm".

At this point, we ought to remind ourselves of the powerful impression made on these philosophers in the middle of the 20th century by the publication of *Philosophical Investigations*. This influence is strengthened by Ryle's eloquent dismantling of a "Cartesian" conception of the mind, advocacy of a reversal of priority between the outer, public, observable doings on the one hand, and the inner, private, and inscrutable mental episodes on the other hand. The effect is a general tendency for philosophers to dismiss the "inner" as explanatorily idle when it comes to understanding and rationality, though their existence was sometimes begrudgingly admitted. Dummett is a prime example of this⁸⁸. It is true that his healthy philosophical instinct lead him to flirt with the idea that linguistic understanding requires "awareness" and "reflection" of some sort, in addition to purely practical abilities to make utterances that fit the occasion. But "awareness" and "reflection" never became a systematic component of Dummett's philosophy of language.

Against this background, it is to Sellars' credit that he labored to develop a "psychological nominalism" and to hold on to the conceptual priority of semantic categories over mental categories *without* entirely abandoning the idea that linguistic actions are expressive of antecedently and independently existing *thoughts*. And by virtue of having a sophisticated philosophy of mind Sellars may turn out to have more resources than Dummett for telling an adequate story about the *ti-esti* of linguistic understanding.⁸⁹

Precisely this virtue is treated as a vice by Sellars' professed student Brandom. His work *Making It Explicit* furnishes us with a very illuminating *negative* example, an example of failure to meet the criterion of adequacy – that our description of

⁸⁸ Another example is Anscombe. Her lecture notes *Intention*, for example, are striking not only for their incisive precision and sensitivity, but also for the conspicuous lack of respect paid to the *mental state* of intending, as well as to the *mental act* of making up one's mind to do something.

⁸⁹ One such resource is Sellars' sophisticated theory of perceptual experience. McDowell, though critical of some of its details, takes it to be amenable to a rationality-grounded-in-experience strategy.

Part I: Semantic Preludes

linguistic practice must make it out to be done with understanding – when one decides to *purge* everything mental from the account of linguistic meaningfulness. To this I now turn.

§2.4.4 Brandom – Mindless Meaningfulness

Brandom’s semantic vision as articulated in *Making It Explicit* can be helpfully⁹⁰ described as an extension of Sellars’ semantic functionalism with two crucial revisions. First, the normative vocabulary, already prominent in Sellars, is now employed “all the way down” in describing the functional system that will turn out to be our linguistic practice. Attendant to this thoroughgoing normativity of functional description is the explicit and systematic adoption of the *pragmatist* understanding of normative elements making up that functional system, something that is only strongly hinted at in Sellars’ work, but never outright formulated and endorsed. Second, Brandom all but severs the intimate connection, maintained by Sellars throughout his semantic writings, between the philosophy of mind and semantics. No mental phenomena, not thoughts, not experiences, are to play any explanatory role in Brandom’s functionalist description of language.

Another departure of Brandom from Sellars is that Brandom’s normative inferential functional description is not meant as an *alternative* to meaning-talk, as Sellars’ is. Brandom makes an *identification* of semantic contents with certain structural features of that functional system. But, like Sellars, being able to play the normative inferential game is supposed to explain what it is to have linguistic understanding. In a later chapter we shall see that Brandom’s theory of semantic content and his theory of linguistic understanding are woven together in the notion of *explicit making*.⁹¹

§2.4.4.1 Functional Description & Semantic Reductionism

After reviewing some existing functionalist approaches to contents of intentional states, Brandom announces that he intends to pursue a “broadly functionalist approach” to content, but “in the context of” linguistic practice and hence to deal with the *semantic* content of expressions rather than psychological states and rational

⁹⁰ Though not completely. Other aspects of the work are being suppressed here, such as its perspectivism and the account of representation based on it.

⁹¹ C.f. §3.4.

agency. In any event, he sees himself as sharing a general explanatory framework with these “reformed” functionalists about psychological states:⁹²

The considerations assembled in the first two chapters [of *Making It Explicit*] suggest the motivation that these two approaches have in common: states, attitudes, and performances are intentionally contentful in virtue of the role they play in *inferentially* articulated, implicitly *normative* practices.

– Brandom, *Making It Explicit*, page 148

So “inferentially articulated normative practices” form the genus to which Brandom’s functional system belongs. The functional system itself is somewhat more complicated than both Sellars’ and Dummett’s, and involves three types of items: statuses (of individual speakers) of entitlement and commitment, attributions to someone (else) of an entitlement or a commitment⁹³, and linguistic performances. Each of these items is indexed to a linguistic expression. So a status of being committed is being committed *to a sentence*⁹⁴, an attribution is an attribution of, say, an entitlement *to a sentence*, and a linguistic performance is the uttering of *a sentence*. Now in the functional system there are *proprieties*, and *only* proprieties, that govern the relations amongst these various items. This is the sense of “norms all the way down”: the fabric of the entire functional system is made up of propriety-relations. For example, the assertion of a sentence *p* makes it *permissible* for a listener to attribute to the speaker a commitment to *p* (i.e. permits an *ascription*), as well as gives the speaker a *responsibility* to give reasons for his assertion in case challenges arise (i.e. makes him responsible for making *further assertions*, given challenges),

⁹² Brandom objects to claims that rational-agency based functionalist theory of content cannot avail itself of the “transitional proprieties” (page 148, *Making It Explicit*) governing beliefs and desires.

⁹³ Brandom identifies a species of this, namely self-attribution, with what he calls the *acknowledgement* (of a commitment). This identification is problematic. But acknowledgement as an attitude distinct from (properly third-person) attribution only becomes important later, in Brandom’s account of the representational dimension of our linguistic practice.

⁹⁴ Of course it is strictly speaking to *what the sentence expresses* that a speaker is committed. And that is also how Brandom puts it (“commitment to a *content*”). But he does it only *proleptically*, in the sense that, at this stage of the functionalist explanation, contents are not yet available. It is precisely the *achievement* of the functionalist explanation that it *makes out* these commitments to be commitments to contents. The situation is analogous to the functionalist explanation of the chess game. The description “rook” or “position” are not yet available in the description of the elements making up a functional system. It is the explanation through functional roles that makes out certain pieces to be *rooks*.

Part I: Semantic Preludes

etc.⁹⁵ These sorts of proprieties are in fact what make these linguistic performances *assertions*.⁹⁶

Other kinds of proprieties make up what Brandom calls the “inferential articulation” of the functional system. These are significant because they “articulate the semantic content” of sentences or sentence-tokens. For example, my ascription to a speaker of a commitment to one sentence A (say “p and q”), if done with entitlement, *permits* me to ascribe to the same speaker commitment to another sentence B (say “p”). In such a case, we can describe the propriety as an *inferential relation between sentences A and B*, specifically, a “*committive inferential relation*” from the sentence A to the sentence B. There are in total four kinds of inferential relations: commitment-preserving inferences, entitlement-preserving inferences, incompatibility relations, and reliability inferences⁹⁷. All these are normative relations governing attributions. The sense in which Brandom has an *inferential* semantic theory is that the functional roles relevant for “articulating” semantic content, according to Brandom, comes from these four sorts of broadly inferential propriety relations. The sense in which Brandom has a *normative* functionalist theory of meaning is that these content-determining relations are, in the final analysis, *proprieties* (governing the making of attributions).

The thoroughgoing normative description of Brandom’s functional system is coupled with what I have been calling the *pragmatist* underpinning of functional roles. What this means is that, not only the *statuses* of commitment and entitlement (which are supposed to be “instituted” by attributions of commitment and entitlement), but also the proprieties that govern the “score-keeping” practice itself (i.e., attributions), are said to be “instituted” by the (second-order) *attitudes of taking* this or that to be a propriety^{98,99}. Normativity and the pragmatist interpretation of it

⁹⁵ For this sort of proprieties, that is, proprieties that individuate the *speech act type* “assertion” rather than *semantic content* of what is asserted, c.f. *Making It Explicit*, page 172-180.

⁹⁶ Contrast this, for example, with a question. The raising of a question does not permit its addressee to ascribe a commitment. Rather, it creates a (prima facie) responsibility on the part of the addressee to say something in response. Brandom himself does not describe in his book other speech act types and the corresponding proprieties, for they are not necessary, according to him, for identifying the functional roles relevant for the explication of semantic content.

⁹⁷ The last of these has special systematic significances for Brandom’s semantic project. See the discussion below in the main text. For historical references, see footnote 103.

⁹⁸ Proprieties’ institution through attitudes is the topic of section V of Chapter 1 in *Making It Explicit*. The view is again summarized at the end of the chapter (page 63): “The fifth point, then, is that one

together constitute the first of the two radicalizations of Sellars I mentioned at the beginning of this section.¹⁰⁰

Brandom says that inferential roles “articulate content”. That the notion of “articulating” in this context is to be understood as *identifying* the content with the inferential roles is clear from numerous passages. A representative passage occurs in his summarizing contrast between what he calls the Kantian conception and the inferential conception of concepts:

The inferential role, which is the conceptual role, *is* the content
– *Making It Explicit*, page 618, (italics Brandom’s original)

This is a typical functionalist move – to identify difficult-to-account-for items with functional roles of sorts – and as such aligns Brandom with his teacher Sellars, even though ultimately, Sellars wants to *replace* the talk about meaning and content with talks about functional roles, rather than analyze the former in terms of the latter. On the other hand, however, Brandom is equally interested in the Wittgensteinian-Dummettian question about understanding.

way to demystify norms is to understand them as *instituted* by the practical attitudes of those who acknowledge them in their practice”. This view – the “retreat from status to attitude” – is taken up again in the concluding chapter, under the name “normative phenomenalism”, in Section II (especially pages 628ff.)

⁹⁹ Brandom also speaks of attributions as well as “acknowledgment” (of a commitment, for instance) as attitudes. But these attributive “attitudes” are explicitly described as items in the functional system. Call them *explicit-attitudes*. These institute *explicit* normative statuses. The attitude of taking certain transitions from one explicit-attitude to another explicit-attitude as *appropriate*, on the other hand, is not officially part of the functional system. They are part of *an additional* account of the constitution of normative relations governing items that *are* in the functional system. In fact, they institute *implicit* proprieties governing the score-keeping practice. In any event, no worry of regress – the worry that the attitudes that institute propriety relations require themselves yet other attitudes – is warranted.

¹⁰⁰ What we have said so far about Brandom’s semantic functionalism is succinctly summarized by himself towards the end of *Making It Explicit* (page 586):

Conceptual content is understood in this work as what can be made explicit in discursive practice. Discursive practice has as its defining core claiming. Claims are a kind of commitment that can be understood in terms of the functional role things of this kind play in social scorekeeping practices – practices and practical attitudes that accordingly can be thought of as instituting this sort of deontic status.

The notion of explicit-making is new, but will become essential in discussions below.

§2.4.4.2 Mindless Meaningfulness

Brandom opens his book on inferential semantics with the declared goal to “make explicit to ourselves who we are”, and immediately begins to answer the imbedded question this way:

What is it we do that is so special? The answer to be explored here – a traditional one, to be sure – is that we are distinguished by capacities that are broadly cognitive. Our transactions with other things, and with each other, in a special way *mean* something to us, they have a *conceptual content* for us, we *understand* them in one way rather than another. ... Picking us out by our capacity for reason and understanding expresses a commitment to take *sapience*, rather than *sentience* as the constellation of characteristics that distinguishes us. ... One of the tasks of this work is to explain what it is to grasp specifically propositional contents, and so to explain who we are as rational or sapient beings.

– Brandom, *Making It Explicit*, pp.4-6

Brandom’s strategy for explaining sapience, or understanding, however, is to identify it as “mastery of proprieties of theoretical and practical *inference*”. We have seen the recalcitrant difficulty Sellars and Dummett ran into when they tried to give a functionalist account that meets the Understanding Criterion. It might be thought that the problem with those accounts was that they took too thin a basis for constructing the functional system. The question now arises as to whether Brandom’s semantic functionalism, which is constructed with thorough normativity built in, might not have more success.

Precisely on the score of rationality, a close cousin of understanding, Brandom’s story has encountered fierce skepticism. It is articulated forcefully by John McDowell in a series of essays. The root cause of Brandom’s failure to deliver an adequate account of rationality and sapience, as McDowell sees it, is the second of what I have identified as two radical modifications of Sellars by Brandom: the total purge of the mental from semantic explanation.

The problem occurs at what I will call the *boundary conditions* of Brandom’s functional system. Specifically, it has to do with the functional role of “observational reports”, which are utterances made with the intention to report what is going on in one’s environment based on one’s observation. The difficulty alleged by McDowell is

that, without a theory of *experience*, a speaker’s linguistic responsiveness to her environment in observational reports can at best be characterized as causal, but not as rational. According to McDowell, it is through experience that a speaker’s report about her environment becomes rationally responsive to it, and thereby intelligible as *significant* and *contentful* speech to begin with, distinguishable from the sound-sequences parrots can be trained to produce. There are two strategies Brandom has employed to save himself from the failure to meet the rationality-of-response requirement. The first is the idea that *grasping* a concept, and therefore understanding, is constituted by one’s capacity to use the concept word *inferentially*, that is, by one’s capacity for certain strictly *intra-linguistic* performances, not by some extra feature of the boundary moves themselves. We might call this the *inferential-imbedding* move. This move addresses the *rationality* of response problem. In fact, the inferential imbedding move is not too different from what Sellars and Dummett already had. Brandom now makes a second move to ensure that observation reports are rational *responses to features of the world*, not just *some* rational behavior. Brandom’s claim is that the thorough-going normative-inferential description does not preclude him from accommodating the *world-directedness* of linguistic practice. The idea is that the reliability requirement – that an observer’s utterance be reliably an indicator of a certain state of affairs in her environment – secures the world-directedness. And the reliability requirement, in turn, is not a causal condition, but is in fact correctly (though “paradoxically”) formulated as a *rule for evaluating purported observational reports*.¹⁰¹ For Brandom, the reliability inference is not made by the speaker, but by the *attributor* who evaluates what the speaker says. Let me call this the *perspective-change* move.¹⁰² It aims to achieve the effect that *all* (content articulating) relations in Brandom’s functional system, including the relation between observational report and extra-linguistic items, are somehow inferential, not

¹⁰¹ C.f. *Making It Explicit*, pp.220-221, pp.188-190. Brandom writes: “Although it sounds paradoxical, for this reason [that is, the fact that observation reports involve the reliability inference on the part of the attributor] the role of a sentence in noninferential reporting should also be understood as falling under the rubric “(broadly)inferential role”. (*Making It Explicit*, page 188-9). We noted already that the reliability-inference associated with observation reports is counted by Brandom as one of four basic forms of “inferential articulation”, next to committive inference, permissive inference, and incompatibility. (Cf. *Ibid.* page 189-90).

¹⁰² This contrasts sharply with Sellars. Sellars has a theory of *experience* to buttress his account of these boundary items, observation reports.

Part I: Semantic Preludes

purely causal.¹⁰³ It is instructive to think of Brandom's "inferentialization" of the reliability condition for observation reports as his replacement for a proper theory of experience underlying observation reports.

But these two moves do not, collectively, get Brandom out of trouble. Let me set aside here the first, inferential embedding move. Suffice it to recall that Dummett's scruple discussed earlier: for linguistic activities to be accompanied by understanding, it may not be enough for the speaker to have the "mere" practical capacities for doing complex linguistic performances; it may be necessary to assume an accompanying *awareness* of some sort. Even if we allow that to pass, and assume that the inferential imbedding move somehow accounts for the presence of *some* understanding, the second move fails to show that this understanding can yield a representational dimension of speech. Inferentialization is a poor substitute for experience.

The perspective-change move fails for a number of reasons. One might focus on the fact that, once we treat the reliability requirement in the form of a rule for attributing commitments, rather than a condition having to do with the interaction between extra-linguistic items with linguistic performances, *no* extra-linguistic items at all will be part of our functional description. That means that what we describe as linguistic practice need not be restrained at all by extra-linguistic states of affairs. But even if we read Brandom charitably and assume that his functional description includes *both* causal (or at least non-normative and non-rational) relations based on reliable responsive dispositions, *as well as* the normative relation from the perspective of an interpreter, the perspective change does not re-instate rationality in a reliable response, when the latter cannot by itself be thought of as a *rational* response.¹⁰⁴

¹⁰³ Reliability condition as an inference is an idea that appears first in Brandom's *Making It Explicit*, relatively late in his writing. It is not in his article *Asserting* or *Varieties of Understanding*. In that article Brandom gives an account of the content-constitutive aspects of a linguistic practice, comprising only the other three kinds of inferential relations. Interestingly, Brandom says in that context that the inferentially articulated content "codifies explicitly" the "merely discriminative classificatory significances implicit in the corresponding non-assertional differential responses". For the idea of expressivism applied to empirical predicates and observational reports, see the next chapter.

¹⁰⁴ McDowell explains the futility of the perspective-change move very clearly in his reply to Brandom's commentary on his book *Mind and World*: "From [the point of view of the putative observer] it [i.e. a putative observation report] appears, in Brandom's conception, not as a rational response to a fact, but as something the putative observer simply finds herself engaged in". This leads to two problems, according to McDowell. First, "[e]ven if we suppose a rational context for the performance can figure in a different perspective, that of an interpreter, it is quite unclear how that could somehow enrich what we can understand to be in view within the perspective of the putative observer, so that the observer's perspective comes to embrace the fact as a potential warrant for reports

Fundamentally, the move does not work for the same reason that the “inferential embedding” move does not work: without rational responsiveness to the environment built directly into the boundary moves such as observational reports, what the speakers have will not be *linguistic* understanding, but at best understanding of some other sort.¹⁰⁵ The “understanding” speakers acquire by virtue of their mastery of intra-linguistic inferential rules, however complicated they are, will be at most akin to the way chess players understand the chess game. But proper *linguistic* understanding must involve the idea of being beholden to, or undertaking to represent, *how things outside the game are*.^{106 107}

In order to characterize the rationality of response in observation reports one *might*, on Brandom’s behalf, point to one of the two attitude-concepts that he mentions to give a pragmatist underpinning of his normative functionalism, namely

or judgments”. Second, and quite obviously, “the putative interpreter would have to be an observer herself. She would have to base her assessments of the performances of the original putative observer on observations, of those performances and their putative subject matter”. Taken together, it “smacks of magic” that, “multiplying what are, considered by themselves, blind responses, to include blind responses to how the blind responses of one’s fellows are related to the circumstances to which they are blind responses, somehow bring[s] it about that the responses are after all not blind”. (*Reply to Commentators*, page 408-9).

I have underlined “blind” because it signals that McDowell’s argument is stronger than it would be if he only said “not rational”. McDowell thinks that the rationality of the restraint on perceptual judgment/reports consists in the restraining fact’s being, as such, “in view” of the subject. See the footnote 105 below on the Williams-McDowell debate.

¹⁰⁵ There is some disagreement as to how such rationality is “directly built into” a boundary move. Michael Williams and Sellars-according-to-Williams (which is probably closer to Sellars himself than McDowell’s Sellars) think that the understanding required is not some additional act or mental occurrence, but a *background recognition*, on the part of the subject, that the circumstance of the observation/perceptual judgment is normal. McDowell thinks that the understanding in question requires a “seeing” (perception), which is conceived as an *occurrence* that makes a rational claim on the subject, rather than an *act* by the subject herself.

¹⁰⁶ This is the point McDowell makes when he writes in §5 of his *Motivating Inferentialism*: “Now as far as I can see, the deontic structure ... that Brandom puts in place in Chapter 3 [of *Making It Explicit*] is consistent with the possibility that a game describable in those terms is just a game, a behavioral repertoire whose moves do not have a significance that points outside the game, so that the moves are not assertions and the transitions are not inferences” (page 127, underlines mine).

¹⁰⁷ In fact, the problem of rational responsiveness is not limited to observation reports, though it is mostly discussed in the literature in that form. Drawing inferences is, like making an observation report, also a *rational* category. To draw an inference is not just to utter two sound-sequences consecutively. To be drawing an inference, the utterance of the consequence must be a *rational response* to the *content* of the antecedent utterance to the extent that, in making the consequence utterance, the speaker must be *aware* that the *ground* for it lies in what the antecedent utterance says. Drawing an inference is therefore not exercising a mere “reliable disposition” to pronounce one sentence after another. It is not clear that Brandom has explanatory resource to characterize the rationality of this kind of response.

Part I: Semantic Preludes

acknowledging (a commitment)¹⁰⁸. But if so, the difficulty would be (a) to understand this attitude of “acknowledging” a commitment, to which Brandom says next to nothing, and (b) doing so might involve characterizing, say, observation reports in ways that are inconsistent with Brandom’s own. For we would then have to describe an observation report as, let us say, a speaker’s “acknowledgement” of a commitment in response to features of her environment, and to understand “acknowledgements” as a rich, and partly psychological category, implicating at least the rationality-grounding role of perceptual *experience*.

§2.5 Conclusions

What the above discussion of Brandom, Dummett, and Sellars hopefully shows, is the sheer difficulty of a *ti-esti* explanation of linguistic meaningfulness, especially when one has denied oneself a theory of the mind. The bottom line is, the *ti-esti* question about such meaningfulness cannot be adequately broached without a) a clearer understanding of the nature of linguistic rules and norms (pragmatist or some other understanding?) and b) an adequate theory of mind to go along with the theory of linguistic meaning. In any event, a *ti-esti* explanation is not what I strive for in this essay. My semantic explanatory ambition here is limited to the *identity question* about meaning.

The explanation of what makes different expressions have different meanings will freely *presuppose* linguistic meaningfulness. Given this limitation, a number of things follow. First of all, the issue of *linguistic understanding* ceases to be the main

¹⁰⁸ “Acknowledging” a status is the only causally efficacious form of the general attitude of *undertaking* a commitment, which Brandom contrasts with the attitude of *attributing* a status. Brandom further claims, implausibly, that undertaking a commitment is merely *self-attributing* that commitment, and for that reason, attribution is the only fundamental attitude (*Making It Explicit*, page 596 ff.).

The reason this is implausible has to do with Brandom’s “perspectival” account – that is, an account based on the perspectival differences between attitudes – of the contrast between facts and appearance. According to that account, what an attributor attributes to another subject shows up for the attributor as an *appearance to that subject*, and what the attributor himself acknowledges shows up for him as *facts*. But if the attributor’s acknowledgement is merely an attribution to *himself*, it is not clear why that does not show up for himself as mere appearance to himself. The problem is not resolved by saying that self-attribution is attribution *de se*, rather than attribution to someone who is *de facto* oneself. For the problem is not that the attributor may, in his *de facto* self-attribution, mistake himself for some other subject. The problem is that, even if the attributor is aware that it is *himself* to whom he makes an attribution, he may still not *stand behind the commitment* he thus attributes. That is, he may still not *acknowledge* the commitment he thus attributes. Witness for example the reluctance with which one often acknowledges the consequence of what one claims, consequence that one clearly *sees*.

focus¹⁰⁹ (much like the issue of rationality recedes from the foreground as soon as one abandons the *ti-esti* explanation of what it is to be human and turns to the corresponding *identity question* instead). We have thereby made the remaining two of the four methodological choices mentioned at the outset of this chapter ((C1) and (C4)). Second, it is no longer a critically important question what *sort of material* out of which we construct our functional system: normative, causal, means-to-end relations or whatever. The reason is *not* because, like for the case of semantic modeling, the only thing about “function roles” of import is their formal characters. The reason is rather that, for *non-constitutional* semantic explanations, we are *not required* to give a description of the function roles that is adequate for the purpose of explaining linguistic understanding, or explaining the nature of linguistic meaningfulness. We can, again, *presuppose* that there is such an adequate description.

Not having to worry about the nature of functional roles in turn relieves our explanatory burden in two respects: it is no longer a pressing task to explain the nature of linguistic *norms* (conceived by a Platonist, a pragmatist, etc.), nor the conceptual relationship between mind and language.¹¹⁰

This concludes our semantic preludes. Enough has been said about general semantic frameworks. Starting with Part II, the focus will narrow, to laws of nature and their linguistic expressions.

¹⁰⁹ This is so notwithstanding the fact that I believe that, and have tried to show why, Dummett is probably right in insisting on the centrality of understanding in any *complete* theory of meaning.

¹¹⁰ Although, depending on the contrast class for the identity question about meaning. Reference to experience might be necessary. Furthermore, other, strictly speaking *non-inferential* roles might be added, the relation of explicit making for example.

Part II Semantic Functionalism & The Problems of Laws

Introduction to Part II

What is the status of law/nomological claims in the tradition of *inferential* (as opposed to possible-worlds) semantics? In §3, I examine the semantic theories about nomological statements or related forms of speech by Dummett, Sellars, and Brandom. Important influences such as Carnap and Ryle are also briefly taken up. As it turns out, most of them cannot or does not maintain the naïve factualism about nomological, or more generally, modal statements: that they state what they appear to state, namely, modal or nomological facts.

Brandom, though strongly influenced by the modal non-factualism of Carnap, Sellars, and Ryle, has a promising rhetoric: *making explicit*. The importance of Brandom's expressivist rhetoric is twofold. First, it suggests an *epistemological* route for empirical information, including nomological information, to reach us. The key idea here is that non-representational *doings* can non-the-less encode some sort of *proto*-information about the world. The second importance of the expressivist rhetoric is that it gives us a richer vocabulary for functionalist semantic explanation: explicit-making in addition to inferences.

Semantic explanation appealing to explicit-making relations, in the case of nomological statements, would appeal to nomological facts/information directly. Such a semantic explanation is obviously not of much use to a *debate* about the very reality of nomological facts, for it presupposes that one side is right. So semantic explanation can only come *after* the epistemological issues are resolved, in §7. On the other hand, the epistemological route suggested by the rhetoric of explicit making can only be followed *indirectly*, in a delicate dialectical maneuver that begins with a *diagnosis and therapy* of skeptical urges. The diagnosis of the root cause of the realism vs. antirealism debate, and a corresponding strategy to nip the skeptical dialectic in the bud, so to speak, requires another idea. This is the idea of *harmony of language use* as a necessary condition for the presumption of contentfulness. This idea I examine and apply diagnostically to laws of nature statements in §4.

§3 Inferentialists on Nomological Necessities & Explicit Making

§3.1 Introduction	109
§3.2 Dummett - Anti-Realism about Abilities	110
§3.3 Sellars and Brandom on Material Rules of Inference	112
§3.3.1 Overview	112
§3.3.2 Constitution Thesis about Material Rules of Inference	114
§3.3.2.1 Sellars	114
§3.3.2.1.1 “Conceptual Meaning”& The Ti-Esti Question.....	115
§3.3.2.1.2 Relational Semantics and “Concept Empiricism”.....	117
§3.3.2.1.3 Myth of the Given & Psychological Nominalism.....	119
§3.3.2.1.4 Forming A Concept vs. Mastering Its Application.....	121
§3.3.2.1.5 “Conceptual” Meaning Explained.....	123
§3.3.2.1.6 One Consequence and One Unresolved Problem.....	124
§3.3.2.2 Brandom Contrasted I – The Meaning-Concept Equivalence.....	125
§3.3.2.3 Brandom Contrasted II – Which Inferences?.....	126
§3.3.2.4 Brandom Contrasted III – Rule-Governed Behavior vs. Rules	127
§3.3.3 Priority Thesis About Material Inferences.....	128
§3.3.3.1 Sellars and Carnap on “Transposed Mode of Speech”	128
§3.3.3.1.1 Carnap.....	128
§3.3.3.1.2 Sellars.....	130
§3.4 Brandom’s Expressivism	135
§3.4.1 Ryle – “Inference Tickets” & “Hypotheticals”	135
§3.4.2 Brandom on “Making Explicit”	139
§3.4.2.1 Overview	140
§3.4.2.2 Observation Reports	142
§3.4.2.2a “Discrimination of Appropriateness” Made Explicit.....	144
§3.4.2.2b “Information” Made Explicit.....	144
§3.4.2.3 Revolutionary Rhetoric.....	146
§3.4.2.4 Logical Vocabulary.....	149
§3.4.2.4a The “Nomological Conditionals”	150
§3.4.2.4b Do Nomological Inferences Have Proto-Content?	151
Appendix to §3.4.2.4	153

§3.1 Introduction

In this chapter, I turn to the three main proponents of inferentialism that we examined earlier to find what they had to say about nomological statements, and whether they have any semantic stories that are compatible with the naïve view that nomological statements describe law facts. It turns out that Dummett says very little in the way of a positive account of the matter at all. He does turn out to be a steadfast proponent of a certain metaphysical anti-realism about ability-statements – a close relative of nomological statements – on semantic ground. Sellars and Brandom have, by contrast, spilled a lot of ink on some very closely related topics. First, there are the so-called *material rules of inference*, which play a central role in their semantic theories. Second, these two authors have similar things to say about the relation between material rules of inference and nomological modal statements. Yet, for both Sellars and Brandom, these do not add up to anything like a direct inferential semantic theory about nomological statements.

What I shall do in this chapter is twofold. First, in §3.2 and §3.3, I will sort out the various elements in the writings of Dummett, but mostly in the writings of Sellars and Carnap, that are relevant to nomological statements. In §3.4 I turn to Brandom, and his notion of “explicit making”, or semantic expressivism. This rhetoric suggests a way to break free of the non-factualism about modal statements found in Sellars and Carnap. I explore what *could* be done to construct a semantic account on the basis of certain interpretation of explicit-making that is compatible with naïve factualism. It turns out that there are problems to this project, some of which can only be solved after Part III.

As a matter of terminology, it will be convenient to speak of “nomological conditionals” and “nomological modal statements”. Conditionals and modal statements are notoriously heterogeneous. Controversies abound as to how to classify them and how the sub-types relate to each other. But it is relatively widely agreed that laws are frequently expressed in conditional, or modal form (“A Y will occur if an X occurs”, “An X is necessarily/followed by a Y”). It is these conditionals and modal statements that are of the most importance to us, and we shall accordingly use the attribute “nomological” to mark them out.

§3.2 Dummett - Anti-Realism about Abilities

The most direct comments from Dummett on issues related to nomological statements come in a *negative* form: he argues *against* what he calls the “naive realism” about statements concerning *abilities*, namely, against the idea that a statement such as ‘John is good at learning languages’ purports to describe a basic kind of reality, which in Dummett’s lingo means that it is capable of being “barely true” and is determinately true or false.¹¹¹ Dummett thinks that either such statements do not describe “some permanent feature of reality” at all, or, if they do, the feature in question is more directly described some other way (e.g. in terms of a *physiological property* of John, to use our example). To put it more succinctly, there are, in Dummett’s view, no basic ability-facts in the world.

Dummett develops this thesis in the context of his assault on the idea that the notion of truth can be made to serve as the basis of a theory of meaning.¹¹² In that context, ability-statements serve for Dummett as an example closely related to one of the three kinds of statements that are “undecidable” and therefore in Dummett’s view present difficulties for the truth-conditional conception of linguistic understanding. Recall that Dummett’s argument has two premises. First, given the truth-conditional conception of meaning, the knowledge of what a sentence means has to be construed as the capacity to recognize whether its truth-condition obtains.¹¹³ Second, certain classes of statements – especially counterfactual conditionals, statements about the past, and quantification over infinite domains – are statements for which there are no procedures allowing one to determine unfailingly whether the condition of their truth obtains. They are in this sense “undecidable”. The first of these premises is tied to Dummett’s verificationism. It is in Dummett’s discussion of the second premise that he argues for his anti-realism for abilities.

The notion of “barely true” is introduced in the context of that discussion. The second premise in Dummett’s argument against truth-conditional semantics, namely the thesis that the three kinds of statements named are undecidable, is argued for on

¹¹¹ For this formulation of “naïve realism”, see Dummett’s *What Is a Theory of Meaning ? (II)*, pages 57ff., as well as the discussion below.

¹¹² See §2.2.1.a, the section titled “*Dummett’s Argument Against Truth-Conditional Conception of Meaning*” for an exposition of the basic structure of his argument.

¹¹³ See §2.2.1.a for discussion of this premise in relation to Dummett’s verificationism.

the basis of the assumption that there are only two ways in which a person can recognize the obtaining of the truth-condition of a statement. The first case is the basic case, where the statement in question belongs to a class whose members are capable of being barely true, which means that an member of that class “cannot be true unless there is some statement, not involving [statements of the class in question], whose truth renders [the first statement] true”.¹¹⁴ The second case is the reductive case, where the statement is not capable of being barely true, in which case in order to grasp its meaning a person must “grasp the way in which its truth depends upon the truth” of some other statements.¹¹⁵ The specific argument that the three classes of statements named are undecidable is based on the further claim that, in the case of statements capable of being barely true, the only model available for truth-condition knowledge is “the capacity to use the sentence to give a report of observation” (page 57).

One can dispute the undecidability premise by, for example, challenging the claim that observation report is the only model for knowing truth-conditions of statements capable of being barely true.¹¹⁶ The other premise of Dummett’s argument against the truth-conditional conception of meaning is problematic as well, as we already noted in §2. For it is not clear why recognition that the truth-condition of a statement obtains is the *only* way one can count as manifesting knowledge that the statement has that truth-condition. My aim here however is not to examine Dummett’s argument against truth-conditional conception of meaning *per se*, but to highlight the fact that Dummett takes its second premise – the undecidability argument about a class of statements – to not only tell against truth-conditional theory of meaning, but to establish an *anti-realist* position about these statements as well. It is one thing to say that there is no procedure that allows one to determine conclusively the truth-value of an ability statement (undecidability thesis about ability-statements). It is altogether another thing to say that the ability statement does not, *therefore*, possess a determinate truth-value (anti-realism about ability statements). Patently, this transition is legitimate only under the additional premise of truth-verificationism: that a

¹¹⁴ Ibid. page 53

¹¹⁵ Ibid. page 57.

¹¹⁶ In fact, this claim exhibits the very presupposition to be identified later in this essay as the main cause of a persistent urge towards a specious skeptical dialectic about laws of nature: that observation is the only basic form of epistemic access we have.

Part II: Semantic Functionalism & Problems of Laws

statement has a determinate truth value only if it is in principle possible for us to conclusively establish that truth value.

Dummett's anti-realist claim about ability-statements therefore depends on two things: an undecidability thesis about these statements, and a version of verificationism about truth. To conclude discussion on Dummett, let me note that the argument from undecidability is by no means the only argument Dummett has for anti-realism *in general*. The most prominent effort by Dummett is contained in his book *Logical Basis of Metaphysics*, which is a sustained argument, on grounds internal to logical systems, against those systems in which the principle of bivalence holds. The silent reliance on truth-verificationism in the context of his critique against truth-conditional conception of meaning, therefore, need not affect the overall strength of Dummett's anti-realist position.

§3.3 Sellars and Brandom on Material Rules of Inference

§3.3.1 Overview

Sellars is responsible for making the concept of “material rules of inference” central to contemporary semantic functionalism. Brandom popularized it. Material rules of inference are close relatives to nomological claims. They are, to put it simply, inference rules whose validity depends on at least some non-logical vocabulary, paradigmatically, empirical predicates. The term “rules” is meant to indicate that *propriety* or *normativity* is a central feature, but leaves open which form that feature takes (e.g. explicit or implicit). A typical example of a material rule of inference is: A is west of B; so, B is east of A. Another type of examples prominent in Sellars' and Brandom's writings are inference rules correlated with laws of nature: if As and Bs are related by a law of nature, then “there is an A; so, there is a B” will be a material rule of inference. Exactly what the relation is between the law and the inference-rule is something that will become the focus of discussion later on.

There are two large theses about the material rules of inference central to Sellars' and Brandom's discussions of them. The first is what I shall call the *constitution thesis* about material rules of inference. It appears in Sellars and Brandom in sufficiently similar form for us to give it a generic formulation:

§3 Inferentialists On Nomological Necessities & Explicit Making

(C – Generic Version) Material rules of inference (partly) constitute the meaning of empirical predicates.

The second, *priority thesis* is about the relation between material rules of inference and nomological modal claims, and there are two variants of them:

(P – Sellars version) nomological modal claims “convey” though do not assert something about the behavior of speakers in connection with certain empirical predicates, namely, that it is in accordance with the corresponding material rules of inference.

(P – Brandom version): nomological modal claims, like conditionals, “make explicit” the material inferential proprieties governing relevant empirical predicates.

Each of these theses will be explained and examined in detail shortly. What makes the last two formulations merit the designation of a *single* thesis rather than two unrelated claims is that Brandom, in formulating his “expressivist” version of it, explicitly acknowledges that he is thereby giving an interpretation to Sellars’ “improvement” of the Carnapian claim of transposed mode of speech. The latter “improvement” on Carnap is just what we are designating as the Sellarsian version of the priority thesis.

The constitution thesis, on the other hand, does not occur without variations either, and we shall see presently what some of these are. The role of this thesis in the semantic thought of Brandom and Sellars is twofold: it has a side of *semantic reductionism* about meanings qua entities, and it has a side connected to the count of *linguistic understanding*. Both sides of the thesis are found in Brandom. For Sellars, the thesis figures primarily in his account of linguistic understanding. (We recall that, with regard to meanings qua entities, Sellars is an *eliminativist*.)

Since I have, in §2, excluded both semantic reductionism as well as an account of linguistic understanding as relevant topics of this essay, the constitution thesis has no direct bearing for the main tasks beginning at §5 below. Nonetheless, I shall begin with a brief examination of the constitution thesis for material rules of inference. It serves as a good testing-ground for some of the methodological distinctions discussed in the last chapter, especially that between the *ti esti* and the *identity* semantic explanation. It also gives us an opportunity to examine in more detail the differences between Brandom and Sellars.

§3.3.2 Constitution Thesis about Material Rules of Inference

§3.3.2.1 Sellars

For all the crassness of Brandom's radicalization and simplification of Sellars' theory of meaning and mind discussed earlier – e.g. the thoroughgoing pragmatist interpretation of norms and complete purge of the mental from a theory of semantic understanding – it is Sellars who holds the more radical, or the “purer” version of the constitution thesis about material rules of inference. Sellars' clearest statements of the thesis are found in a couple of papers from the early 50s:

Stated summarily [the position I wish to defend] claims that all conceptual meaning, the conceptual meaning of descriptive as well as logical symbols, is constituted, completely constituted, by syntactic rules. ...

Let me now put my thesis by saying that the conceptual meaning of a descriptive term is constituted by what can be inferred from it in accordance with the logical and extra-logical [e.g. material] rules of inference of the language (conceptual frame) to which it belongs.

– *Is There A Synthetic A Priori?* page 136 (underline mine)

... there is an important difference between logical, modal and normative predicates, on the one hand, and such predicates as “red” on the other. In the case of the former, it is obvious that their conceptual meaning is entirely constituted by their “logical grammar”, that is, by the fact that they are used in accordance with certain syntactical rules. In the case of the latter, this is not obvious – though, as we are about to argue, it is equally true.

– *Inference and Meaning*, page 282/334. (underline mine)

If ... the argument of section V is sound, it is the first ... alternative to which we are committed. According to it, material transformation rules determine the descriptive meaning of the expressions of a language within the framework established by its logical transformation rules. The familiar notion (Kantian in its origin, but present in various disguises in many contemporary systems) that the form of a concept is determined by ‘logical rules’, while the content is ‘derived from experience’ embodies a radical misinterpretation of the manner in which the ‘manifold of sense’ contributes to the shaping of the conceptual apparatus ‘applied’ to the manifold in the process of cognition.

– *ibid.* pages 284-5/336-7 (underline mine)

A striking feature of Sellars' formulation is his insistence on the attribute "conceptual" whenever he speaks of meaning in these texts. A related peculiarity is Sellars' opposing of "conceptual meaning" of an observational predicate to its "application", even though he takes both to be conveyed¹¹⁷ by a typical meaning statement. Another related feature characteristic of Sellars is that the latter, the application of a word, is a matter of "conditioning and association" alone, and *not* genuinely rule-governed behavior (though they are governed by *ought-to-be* norms). This set of features can be brought under the question: what does Sellars mean by "*conceptual* meaning" rather than simply "meaning"?

A second characteristic of Sellars' statement of the constitution thesis that puts it in contrast to Brandom's, for example, is Sellars' emphasis on ruled-governed *behavior*, in contrast to abstract rules *per se*. I begin with the first of these two characteristics.

§3.3.2.1.1 "*Conceptual Meaning*" & *The Ti-Esti Question*

Sellars' insistent talk of "*conceptual* meaning" first seems to suggest that he is mainly interested in a certain *conceptual character* of linguistic meaning. Whatever that is, it does not encompass such facts as, for example, competent German speakers have a "conditioned habit" to respond with the word 'rot'¹¹⁸ to "sensory states induced by environmental stimulation – in this case, by red objects"¹¹⁹. However, it would clearly be wrong to characterize Sellars' view by saying that the latter kind of facts – of "conditioning or association" – has nothing to do with meaning *per se*. Just the opposite, Sellars writes:

Might it not be the case that Smith's statement "When [Fritz]¹²⁰ says 'rot' it means red" (where 'red' is granted to be an observation predicate) conveys

¹¹⁷ For the distinction "convey" vs. "assert", see the discussion on Sellars' version of the priority thesis below, in §3.3.3.1.2 Sellars.

¹¹⁸ I have systematically standardized Sellars examples to be about German expressions.

¹¹⁹ Is There A Synthetic A Priori? page 134-5.

¹²⁰ In accordance with the systematic alteration of Sellars' examples, I have changed 'Jones' to a more common German name.

Part II: Semantic Functionalism & Problems of Laws

- (1) [Fritz] has habits with respect to ‘rot’ which we could correctly speak of in terms of *rules* for using ‘rot’ over and above logical rules in the narrow sense,
- (2) [Fritz] (roughly) is conditioned to respond to red objects with ‘rot’?

This is the thesis I wish to defend ...

– *Is There A Synthetic A Priori?* page 135.

We will come to the business about “convey” vs. “assert” later when we discuss the priority thesis. Sellars’ view is that meaning statements do not directly *assert* either (1) or (2). But this aspect of Sellars’ thinking does not undermine his view expressed by the passage cited and many others, which is that conditioning by and association with extra-linguistic entities (i.e. (2)) is *just as relevant* to meaning as the “habits” of use in accordance with logical and extra-logical rules (i.e. (1)) are. Sellars himself never officially clarifies what he means with the attribute “conceptual” in this semantic context. But it is, as I shall argue presently, a plausible paraphrase to say that, in zooming in on “*conceptual* meaning”, Sellars is pursuing his interest in the very character of *meaningfulness* in general (“in virtue of what is E a meaningful expression at all?”) and the related question about the nature of linguistic understanding in particular, rather than in the sort of semantic explanation we have identified as the “identity” question about meanings (“in virtue of what does E mean what it means?”).

It is easy to see why Sellars’ focus *should* not be on the identity question. For if Sellars’ main goal was to explain why ‘rot’ means red whereas ‘blau’ means blue, he could have pointed to the “conditioning and association” component (2) alone for a satisfactory answer: ‘rot’ means red rather than blue because competent German speakers are conditioned to respond with it to *red* things, not to blue things. Textually, the most telling sign for what Sellars is really up to is his – historically extremely interesting – discussion and critique of what he calls “concept empiricism” in his paper *Is There A Synthetic A Priori?*. It shows that Sellars’ semantic thinking is inextricably tied up with his thinking on (broadly) epistemological issues that form a major topic of his celebrated *Empiricism and the Philosophy of Mind*. The lectures that were to become the basis of that book took place three to four short years after Sellars’ early papers on semantics in the early fifties. In particular, elements of the discussion of “concept empiricism” are taken up and expanded in *EPM*, specifically in sections VI-VIII (§26-38). These thematic correlations make it clear that, among

other things, Sellars was already in these early papers zooming in on versions of what he would later call “myth of the given”. These claims need more detailed exegetical exposition.

§3.3.2.1.2 *Relational Semantics and “Concept Empiricism”*

The main thesis in *Is There A Synthetic A Priori?* is that there are statements that are true in virtue of the meaning of the non-logical expressions occurring in them alone.¹²¹ In particular, Sellars wants to entitle himself to the claim that a statement of the form “All As are B” is true in virtue of the meanings of empirical predicates “A” and “B”, when there is a material rule of inference from “x is A” to “x is B”. To do that, Sellars attempts a paradigm change in thinking about meanings. The initial conception with which Sellars starts is a *relational* one. It holds that an empirical predicate has a universal or a property as its “*real meaning*”¹²², and the fact of its meaning such and such consists in its being related to its “*real meaning*”. Now, for there to be a synthetic *a priori* in Sellars’ sense in the form of “All As are B”, there must be a “*real connection*” between the real meanings of the two empirical predicates “A” and “B”, or in other words, a “synthetic necessary connection” between two properties.¹²³ At this point, the existence of such synthetic *a priori* statements appear extremely doubtful against the background of the so-called *concept empiricism*, a broad view about concept acquisition which says:

Concepts of qualities and relations are formed from particulars [that exemplify these qualities and relations]”.

– Ibid. Page 128.

The reason that concept empiricism is a problem for Sellars is that it seems to exclude the very possibility of forming the concept of “real connections”:

The implication of concept empiricism with respect to the concept of real connection is immediate and murderous. There is no such concept. (...) The truth (...) is that if there is such a thing as necessary connection, it is a relation

¹²¹ Truth in virtue of meanings alone, or as Sellars puts it, *ex vi terminorum*, is one way to spell out what Sellars means by “a priori”.

¹²² That turn of phrase is contrasted with “linguistic meaning”.

¹²³ *Is There A Synthetic A Priori?*. Page 127-8.

Part II: Semantic Functionalism & Problems of Laws

satisfied by *universals* (a relation whose terms are universals), and *not* by particulars. Thus, for the concept empiricist, our failure to have such a concept is not a mere matter of failing to find any particulars which exemplify it; we couldn't find particulars which exemplify it.

– *Is There A Synthetic A Priori?* page 129.

[In this short passage, Sellars inadvertently lays bare the forces that moved Kant to his influential modal anti-realism: the concept of necessary connection is satisfied by universals, but we cannot get *empirical warrant* for the instantiation of the necessary-connection relation by two universals. I spell out this in the form of an argument against Kant in §6.2.2] The problem, then, results from the combination of a certain form of *relational conception of meaning* that entails that synthetic *a priori* truths in Sellars' sense require a kind of relational concepts exemplified by universals on the one hand, with a certain empiricist conception of *concept formation* that denies the very possibility of forming such concepts, on the other hand. (As we shall see presently, (some versions of) “concept empiricism” may allow concepts of logic and mathematics to be exceptions to the supposed necessary mediation by particulars for concept formation. The concept of “real connections” however is unlikely to have that exceptional status, given that it is supposed to be a concept of something *empirical*, of an aspect of reality independent of the inner-workings of our language.)

Now given that this combination of a semantic view with a (broadly speaking) epistemological outlook precludes the acknowledgement of synthetic *a priori* truths in Sellars' sense, what must be rejected and/or modified? Note first that, if the relational semantic view – e.g. that the fact of ‘red’ meaning *red* consists of ‘red’s relation to its “real meaning”, the property of redness – is kept, there is not much wiggle room for Sellars on the side of epistemology, especially given the following highly plausible equivalence:

(Meaning-Concept Equivalence):

To grasp the concept *red* is to grasp the meaning of “red”.

For, the relational semantic view regards phrases like “the meaning of E” *categorically*, as a referential singular term with an abstract entity as referent, and sentences like “E means ...” as expressing a relation between an expression and an abstract entity. Under this view, to grasp the concept *red* is to grasp the meaning of

“red” (by meaning-concept equivalence), which is in turn to stand in a grasping relation to a universal (by categorematicity of “the meaning of ‘red’”). Now how does one come into this grasping relation to a universal? Concept empiricism says by getting in contact with particulars exemplifying that universal. What other option is there? Given that the topic is empirical predicates, the rationalist option of mentioning the faculty of “intellectual intuition”, say, is simply out of the question.

What this brief discussion shows is that the relational semantic view is not a thesis unrelated to “concept empiricism”: it makes the latter practically the only plausible alternative. Moreover, we shall see presently that one major form of “concept empiricism” will be rejected by Sellars as the “very heart of the Myth of the Given”. Against the background of a tight connection between the relational semantic view and “concept empiricism”, it makes sense to see the relational semantic view itself as the semantic counterpart to Myth of the Given. A workable formulation of this perspective is:

(Relational Semantic view as Semantic Myth of the Given):

There are, independent of and antecedent to any linguistic practice, “real meanings” in the world; and for a linguistic expression to be meaningful is for it to match onto one of these antecedently available “real meanings”.

The combination of views that preclude synthetic *a priori* statements in Sellars’ sense, therefore, is the semantic Myth of the Given combined with “concept empiricism” which is practically the only kind of views about concept formation compatible with the semantic Myth.

§3.3.2.1.3 *Myth of the Given & Psychological Nominalism*

The version of concept empiricism that I said anticipates the official, that is, epistemological, “heart of the Myth of the Given” in *EPM* is the first, “mental eye” variety:

Thus the concept empiricist of this brand conceives of such symbols as “red” and “between” as acquiring meaning by virtue of becoming associated with such abstract entities as redness and between-ness, the association being mediated by awareness of these entities.

– Ibid. Page 130 (Underlines mine).

Part II: Semantic Functionalism & Problems of Laws

To see why that is just what Sellars later in *EPM* calls the “Myth of the Given”, one only has to take a look at the example Sellars gives for this sort of concept empiricism, namely, C.I. Lewis:

I am afraid, however, that our agreement with Lewis is more shadow than substance. For while he writes in this manner of the interpretation of the given by means of concepts whose implications transcend the given, he also holds that the sensible¹²⁴ appearances of things *do* wear their hearts on their sleeves, and that we do have a cognitive vision of these hearts which is direct, unlearned and incapable of error – though we may make a slip in the expressive language by which these insights are properly formulated.

– Ibid. Page 130 (Underlines mine).

Having read this passage in *Is There A Synthetic A Priori?*, it is hard to miss what Sellars has in mind in pronouncing the following indictment in the *EPM*:

The idea that observation "strictly and properly so-called" is constituted by certain self-authenticating nonverbal episodes, the authority of which is transmitted to verbal and quasi-verbal performances when these performances are made “in conformity with the semantical rules of the language,” is, of course, the heart of the Myth of the Given, ...

– *Empiricism and the Philosophy of Mind*, §38.

There is more to the correspondence between *EPM* and the early paper on synthetic *a priori*. In those passages of *EPM* surrounding the passage just quoted, Sellars opposes the idea of “self-authenticating”, unlearned, incorrigible awareness of qualities – the “very heart of the Myth of the Given” – with what he calls *psychological nominalism*:

If, however, the association [between words and “classes of resembling particulars”] is not mediated by the awareness of facts either of the form *x resembles y*, or of the form *x is f*, then we have a view of the general type which I will call *psychological nominalism*, according to which *all* awareness of *sorts, resemblances, facts, etc.*, in short, all awareness of abstract entities – indeed, all awareness even of particulars – is a linguistic affair.

¹²⁴ The word is used here with the sense of “capable of being taken up by the *senses*”.

§3 Inferentialists On Nomological Necessities & Explicit Making

– Ibid. §29. (Underlines Mine)

What Sellars calls *psychological nominalism* in *EPM* is a “view of [a] general type”, of which Sellars’ own view will be a species. It in fact includes as a sub-type the *other* version of concept empiricism discussed by Sellars in the earlier essay:

In its traditional form, this second approach ... insists that this association develops by the joint occurrence in the mind of instances of the word and of the characteristic in question, in this case redness, unmediated by awareness of abstract entities. In other words, while it is redness that is associated with “red,” the mechanism whereby this association is created does not involve awareness of redness, but only the joint occurrence in experience of instances of redness with tokens of “red”.

– *Is There A Synthetic A Priori?* page 132. (my underlines)

The genericness of “psychological nominalism” lies in the sweeping turn of phrase “linguistic affair”. The second form of “concept empiricism” specifies the relevant linguistic affair to be joint occurrences of red objects with tokens of the word “red”. This species of “psychological nominalism”, however, is not the one ultimately endorsed by Sellars.

§3.3.2.1.4 *Forming A Concept vs. Mastering Its Application*

Sellars develops his positive, non-relational account of meaning statements, an account we have already encountered, by first working out a view of concept-formation obtained through a series of modifications of the version of “concept empiricism” falling under “psychological nominalism”.

First, Sellars observes that the “joint occurrence”/association of words with objects alone will be inadequate for concepts of logic and mathematics. This leads to a “usual” remedy of the theory by recognizing “a second mode of concept formation, namely the learning to use symbols in accordance with rules of logical syntax” (i.e. logical rules of inference). Then, Sellars remarks that “even those terms, such as ‘red,’ which are supposed by the theory to gain meaning by association, share in the second mode of concept formation, for only by being used in accordance with rules of logical syntax can they perform the functions by virtue of which a concept is a concept” (Ibid. page 133. Underlines mine).

Part II: Semantic Functionalism & Problems of Laws

From these passages one might reconstruct the following line of thought. On the one hand, the case of logical and mathematical concepts shows that the “joint occurrence” of linguistic tokens and extra-linguistic entities is, in some cases, *not necessary or even relevant at all* for linguistic expressions’ expressing concepts. The case of empirical predicates and examples like thermometer and well-trained parrots, on the other hand, shows that such reliable joint occurrence is in all cases *not sufficient* for words’ expressing concepts. The idea is crystallized in Sellars’ repeated emphasis of the point that responding to red objects with tokens of ‘red’ cannot be a matter of following rules. It is spelled out a bit more in an essay written around the same time:

... the fact that both rule-governed and merely associative uniformities [of linguistic behavior] are *learned* uniformities, and differ in this respect from, say, the uniformities studied in chemistry, has blinded many philosophers to the important respects in which they differ from one another ...

– *Inference and Meaning*, page 284/336.

The point is that the “associative uniformity” of one’s conditioned response to red objects with tokens of ‘red’ is indeed learned, but it is in the same sense of “learned” as when we say that a parrot can *learn* to greet his master with tokens of “good morning!”. But the latter clearly *has nothing to do* with mastering a concept.

Considerations such as these lead Sellars to regard the learning to use ‘red’ in accordance with relevant rules of inference a matter of “mastering the concept” proper, even though it does not exhaust *all* the use of that word. What one learns *in addition* to the rules, namely, to reliably respond to red objects with tokens of ‘red’, is not part of mastering the *concept* of red, but to master its *application*:

... a language or conceptual system has two dimensions, an intra-systemic dimension in which the elements of the system are related to one another by syntactical rules; and an extra-systemic dimension in which the system gains *application* by virtue of an association of some of its terms with extra-conceptual reality. A conceptual frame is a rule-governed system. Its application is a matter of association rather than rules.

– *Is There A Synthetic A Priori?* page 133.

§3.3.2.1.5 “Conceptual” Meaning Explained

We are now finally in a position to appreciate the cause of Sellars’ curious turn of phrase “conceptual meaning”. Earlier we saw that Sellars takes meaning statements to convey two sets of facts: an expression’s being used in accordance with certain rules of inference, and its use in conditioned response to extra-linguistic entities. The way Sellars came to this semantic view is by way of endorsing (in effect, his own version of the *Priority of Understanding Thesis*¹²⁵) the idea that *meaning* statements convey exactly the same sets of facts that are constitutive of the *mastery* of a concept and of its application, respectively. I shall call it the distinctively Sellarsian correspondence of (concept)-epistemology and semantics, or more briefly, the Meaning-Mastery Correspondence. Here is a tabular representation of it:

<i>Conveyed by Meaning Statement:</i>	<i>Mastery of Concept/Application:</i>
‘red’ <u>is used</u> by competent speakers in accordance with material rules of inference;	S learns to <u>use</u> ‘red’ in accordance with material rules of inference;
⇔	
‘red’ <u>is used</u> by competent speakers as conditioned response to red objects;	S learns/is conditioned to <u>respond</u> to red objects with tokens of ‘red’;
⇔	

For Sellars, the (concept-)epistemological side is, at least methodologically, prior of the two (hence a version of his “Priority of Understanding Thesis”). And on this side, Sellars has adopted the view that only the *first* of the items in the “mastery” column is properly related to *concept* formation, or, we might say, *linguistic understanding*. On the semantic side, he wants to maintain the intuitive, and philosophically sane, view that both items in the “conveyed by meaning statement” column pertain to meaning. The correlation between the epistemological and semantic side then leads (or forces, depending on one’s view) him to describe the use of an expression according to rules of inference as constituting its *conceptual* meaning. Now, if we had to give a name to the second factor conveyed by a meaning statement, we would have to call it the “*applicative* meaning” of the expression in question. Incorporating Sellars’ consideration on rule-governedness, we can represent the meaning-mastery correspondence in a more schematic manner:

¹²⁵ C.f. discussions of Dummett in the previous chapter, especially §2.2.1b.

Part II: Semantic Functionalism & Problems of Laws

	<i>Meaning</i>	<i>Mastery</i>
rule-governed:	“conceptual” meaning of ‘red’	⇔ mastery of the concept <i>red</i>
(merely) learned/conditioned:	“applicative” meaning of ‘red’	⇔ mastery of its application

The reference to “*conceptual* meaning” in the formulation of Sellars’ constitution thesis about material rules of inference is therefore a reflection, through the meaning-mastery correspondence, of Sellars’ view that conditioned response to extra-linguistic entities is not part of mastering an empirical *concept*. That is, it does not figure in an account of what it is to *understand* the expression ‘red’. It also ensures that Sellars’ semantic view does not amount to a straightforward *linguistic idealism*, which would be the case if Sellars had maintained the constitution thesis for “meaning” *tout court*.

§3.3.2.1.6 One Consequence and One Unresolved Problem

Note that in the picture Sellars arrived at, the semantic Myth of the Given has become a fundamental impossibility. For usage in accordance with inference rules, what is supposed to constitute the *conceptual* side of meaning facts, is something that is simply unintelligible independent of and antecedent to linguistic practices.

On the other hand, there are some not necessarily desirable consequences. Note first that Sellars’ meaning-mastery correspondence coupled with his view on concept formation brings him into direct conflict with what we have called the “meaning-concept equivalence”. For, if to grasp the concept *red* is only a matter of learning to use ‘red’ in accordance with certain rules of inference, and the latter is only *part* of what is relevant to meaning, then obviously grasping the concept of *red* will not be equivalent to, but rather strictly less demanding than grasping the *meaning* of ‘red’. Sellars can therefore only endorse a modified version of we have called the meaning-concept equivalence, where “meaning” is replaced by “conceptual meaning”:

(Meaning-Concept Equivalence – Sellars Version):

To grasp the concept *red* is to grasp the conceptual meaning of “red”.

This formulation seems to conflict with our intuitions about how meaning and concept are related.

More relevant to the topic of this essay, however, is a problem that remains unresolved by Sellars' philosophical maneuvers. For, inquisitive readers of Sellars will point out that the concept of "real connections", which is what gave Sellars' trouble in his pursuit of the synthetic *a priori* to begin with, remains problematic. Granted that there are rules of use such that, grasping them constitutes grasping the *concept* of "real connections", one still wants to know what constitutes mastering the *application* of this concept. The application of *this* concept cannot, as Sellars himself observed, be mediated by particulars exemplifying it, for there aren't any such particulars. So mastering its application cannot be the same as mastering the application of an empirical concepts, that is, a matter of being conditioned to respond to particulars with tokens of a word particulars falling under that concept. In fact, Sellars does not provide an alternative account of concept application for the case of necessary "real" connections.¹²⁶ Another option would be to assimilate the concept of "real" connections to those of logic and mathematics, and say that it, like the latter, does not have any empirical applications in Sellars' sense. Sellars' silence on this matter, together with Sellars' tendency to deny, or at least be quiet about, the "descriptive function" of nomological necessity claims, makes a plausible suspicion that Sellars' priority thesis about material inferences (over nomological statements) leaves him little room besides the anti-realist position which he himself would more than hesitate to explicitly endorse. We shall return to this below.

§3.3.2.2 *Brandom Contrasted I – The Meaning-Concept Equivalence*

To summarize the discussion of so far, let me give the following more precise formulation of Sellars' constitution thesis about material rules of inference:

(C – Sellars Version) The *conceptual* part of what a meaning statement about an empirical predicate *conveys* is constituted by the use of that predicate in accordance with material rules of inference.

¹²⁶ That Sellars thinks there *exists* such a concept seems beyond doubt given his elaborate handling of it in an essay frequently cited by Brandom: *Concepts as Involving Laws and Inconceivable Without Them*.

Part II: Semantic Functionalism & Problems of Laws

We saw that the qualification of “conceptual” is to signal the importance of “conceptual meaning” to the question of concept mastery, or *understanding* of the predicate in question.

How does Brandom differ from Sellars in the constitution thesis, which he has inherited from Sellars? There is the methodological difference already mentioned: Sellars intends the Constitution Thesis to be primarily part of his account of concept mastery and linguistic understanding, whereas Brandom also intends it to be part of his semantic reductionism. There are also some substantive differences.

One substantive difference is that Brandom’s version of the constitution thesis does not violate the meaning-concept equivalence. Brandom maintains the same meaning-mastery correspondence of Sellars; but he does not draw the distinction between mastering a concept vs. mastering its application. Both items in the “mastery” column are, for Brandom, elements of mastering the concept of *red*. Accordingly, on the side of semantics, Brandom drops the more narrow concept of “conceptual meaning” entirely. So his constitution thesis concerns *both* elements in the “meaning” column. Yet he also maintains the Sellarsian thesis that conditioned response to red objects with tokens of ‘red’ is not rule-governed behavior. So how can this “associative” element be a matter of inference *rules*? Brandom’s solution, as we have already seen in the previous chapter, is to replace the second element in the “meaning” column – the conditioned responses in observation reports – with proprieties for *reliability inferences* governing the *attribution* of a conditioned response to someone making an observational report.¹²⁷ Reliability inferences are not material inferences, as they are inferences from the occurrence of an observational utterance to features in that part of the utterer’s immediate environment which is within his perceptual reach. To that extent, Brandom’s constitution thesis is not *purely* about material rules of inference.¹²⁸

§3.3.2.3 Brandom Contrasted II – Which Inferences?

Brandom’s view is “impure” in another, methodological way. For he leaves open an *entirely different* way of accounting for the meaning of empirical predicates, one that

¹²⁷ C.f. Brandom’s *Making It Explicit*, pages 188-190.

¹²⁸ Sellars has, however, a theory of *experience* for the “applicative” meaning of an empirical predicate. Roughly, Brandom has rules at the boundary of his functional system, Sellars has experiences. The latter option is much more likely to lead to success. For more on this contrast, see the discussions in §2.4.3 and §2.4.4 above.

is palatable only to those convinced of Quine's semantic holism. This alternative option consists of, in Brandom's words, "[putting] *all* the inferences in the same basket, by rejecting the distinction between those good in virtue of meaning and those good in virtue of matters of fact" (*Inferentialism and Some of Its Challenges*, page 662). I shall not pursue this Quinean inferentialism any further in this essay.

§3.3.2.4 Brandom Contrasted III – Rule-Governed Behavior vs. Rules

What meaning statements convey, according to Sellars, is that certain linguistic *behavior* is either "in accordance with rules"/"rule-governed" or conditioned. It is true that he sometimes formulates the constitution thesis loosely by saying that meaning is constituted by material *rules* of inference.¹²⁹ But in more detailed formulations, Sellars always comes back to rule-governed *behavior*.¹³⁰ Moreover, the briefer formulation *cannot* be literally interpreted for the simple reason that a main Sellarsian doctrine says that meaning statements do not refer to "meanings" at all. So Sellars cannot be taken as claiming *anything* about meaning as an abstract entity, let alone identifying it with a set of rules.

Pinning down Brandom on the issue of whether it is rule-governed *behavior* or the *rules* themselves that constitute meaning facts is much harder. It suffices to note that Brandom often speaks of "inferential proprieties" as articulating the semantic content of expressions governed by them.¹³¹ Rather than a deep difference between Sellars and Brandom, this may reflect a general difference in level of willingness to directly talk about abstract entities such as meanings/contents and rules. After all, Brandom's pragmatist "expressivism" – more on that later – takes knowing-how and concrete doings to be more basic than codified rules.¹³² Sellars' reluctance to directly talk about abstract entities is a reflection, as I have indicated earlier, of the nominalist/behaviorist climate of the times in which he wrote. One of its central manifestations is Sellars' across-the-board application of the *convey vs. assert* strategy to analyze claims that are putatively descriptive of various sorts of abstract

¹²⁹ e.g. the first of the three passages quoted above at the beginning of our discussion on this thesis.

¹³⁰ e.g. the second the three passages quoted at the beginning.

¹³¹ For example, "contents articulated in terms of these [inferential] proprieties" (*Making It Explicit*, page 190) is such a turn of phrase that Brandom is fond of.

¹³² One might say that the difference is while Sellars avoids explicitly endorsing the ontological reality of rules and meanings, Brandom insists that they are, though real, secondary to "practices".

entities in terms of their *conveyance* of facts about non-abstract entities, while remaining deafeningly silent about what these claims actually *assert*. We shall see this in action presently, as we move to the priority thesis about material rules of inference.

§3.3.3 Priority Thesis About Material Inferences

The attribute “conceptual” in the constitution-thesis of Sellars has been dissected in the last section, now it is time to turn the limelight on the curious verb “convey”. As already mentioned, Sellars makes the distinction between “convey” and “assert” for various sorts of sentences. What is of interest to us is the distinction as applied to nomological necessity claims, that is, Sellars’ version of the priority thesis about material inferences. For, while the constitution thesis is an especially salient strand in inferentialists’ semantics that serves for us as an important *background*, the priority thesis constitutes the most important thing inferentialists have had to say about our topic *proper*: nomological modality.

§3.3.3.1 Sellars and Carnap on “Transposed Mode of Speech”

Sellars’ distinction of convey vs. assert has its origin in Carnap’s concept of “transposed mode of speech”, but is not identical with it (Brandom’s careless assimilation¹³³ notwithstanding).

§3.3.3.1.1 Carnap

Carnap’s idea is that, many statements are, given their grammatical form, purportedly about non-linguistic objects of some sort, or state non-linguistic facts of some kind, while what they *really state* are facts of syntax or linguistic use. These statements are said to be *quasi-syntactic* statements in the *material* mode of speech (“material mode” because of their apparent *non-linguistic* content). Material mode of speech is only one kind of transposed mode of speech. To be in a transposed mode of speech a statement only need to have an *apparent* topic that diverges from its *real* topic. Metaphor, for example, counts for Carnap as a transposed mode of speech.¹³⁴ Material mode of

¹³³ For example, Brandom writes in the fourth of his *John Locke Lectures*: “The large claim in the vicinity ... is, as Sellars puts it, that ‘the language of modality is...a “transposed” language of norms.’” (page 12). Brandom is quoting a sentence from Sellars’ *Inference and Meaning*. But he apparently does not see that with the sentence Sellars summarizes a position of Carnap, a position that Sellars then proceeds to criticize and modify. The key difference here is that Sellars wants modal statements to convey facts about *behavior*, not *norms* themselves.

¹³⁴ C.f. Carnap’s *Logische Syntax der Sprache*, §80.

speech has the additional character that its *real* content can be stated by another statements that is explicit about *syntax*. The latter is the corresponding statement in the *formal* mode of speech.

Carnap counts incredibly many kinds of statements as being in the material mode of speech. Examples are: statements about numbers, *oratio oblique*, statements in which words like “property”, “object”, “relation”, “state of affairs” are used, and even statements about colors and time.¹³⁵ The example of interest to us is, of course, statements of material (non-logical) necessity. The *real* topics of these statements, Carnap tells us, are linguistic expressions: number *signs*, *sentences* used to relate an assertion, and so on. For Carnap, the notion of material mode of speech serves a broad range of goals: to respond to C.I. Lewis’ charge that intensional expressive resources are needed to express implication (“we don’t need intentional object-language expressions, if we have meta-language means to talk about non-logical transformation rules”¹³⁶); to Wittgenstein’s idea that many categorial facts can only be *shown*, not be said (“of course they can be said, but only in our syntactic meta-language”¹³⁷); and finally, as an approach to almost *any* metaphysical dispute whatsoever (“it’s not about *numbers* that we disagree, it’s about *numerals!*”¹³⁸).

To focus on the case that interests us, we can characterize Carnap’s idea about material necessity claims in terms of two theses. Let X be the material necessity claim, Y be its corresponding claim about rules of inference. For now, we can adapt Carnap’s terminology to that of Sellars (as he is closer to our topic) and set X to be “ ϕa necessitates ψa ” and Y to be “ ψa ’ is derivable from ‘ ϕa ’”. Carnap’s two theses are:

(Sameness of Content)¹³⁹

The statement X says the *same thing* as does the statement about rules of inference Y;

(Non-Literalism about Necessity Statements)

¹³⁵ Ibid. §74-9.

¹³⁶ Ibid. §69-71.

¹³⁷ Ibid. §73.

¹³⁸ Ibid. §74-9

¹³⁹ For Carnap’s formulation of this, see, for example, *Logische Syntax der Sprache* §74.

Part II: Semantic Functionalism & Problems of Laws

The statement about material necessity X, should *not* be interpreted *literally*, but rather by reference to the statement about inference Y.

To *interpret* a statement means here to formulate, in words more familiar or easier to understand, what that statement asserts (as opposed to what it insinuates, suggest, implies, or presupposes). Note that the same-content-thesis leaves all three options open: (i) X and Y are *both* be interpreted literally, for facts about necessity *are* facts about inference rules; (ii) only X is interpreted literally, and Y is interpreted by reference to X; (iii) only Y is interpreted literally, and X is interpreted by reference to Y. Non-literalism about necessity rules out (i) and (ii). The bud that is to blossom into the *priority* of material inferences over nomological necessity claims in Sellars, is none other than this non-literalism thesis.

§3.3.3.1.2 Sellars

These two theses unpack the claim that “the language of modality is ... a ‘transposed’ language of norms”. That formulation is Sellars’ summary of Carnap’s position, not his own.¹⁴⁰ Though in many respects, Sellars maintains these two theses in a somewhat weaker form.

The first deviation from Carnap to note is that, in all Sellars’ official formulations, the inference side of the two relata of Carnap’s necessity-inference relation does not pertain to inference rules *per se*, but linguistic *behavior* in accordance with inference rules. Another deviation from Carnap is that Sellars adamantly renounces the same-content thesis, because, as he plausibly says in many places, the language of norms, language of psychology, and the language of modality cannot be reduced to each other.¹⁴¹ The reason that the language of psychology is of relevance here is that Sellars takes “rule-governed behavior” to be psychological facts

¹⁴⁰ Brandom attributes it carelessly to Sellars in, for example, the fourth of his *John Locke Lectures*. C.f. footnote 133 above for details.

¹⁴¹ For example, “In short, modal terms, normative terms and psychological terms are mutually irreducible” (page 281/333, *Inference and Meaning*). Sellars’ anti-reductionism is quite general. He argues for it in the case of physical, mental, and semantic vocabulary in *Mind, Meaning, and Behavior*. In a footnote of his essay on synthetic *a priori*, Sellars writes “‘descriptive discourse’, ‘discourse about rules’ and ‘discourse about meaning’ are three distinct ‘modes of speech’”. (Foot-note 7, *Is There A Synthetic A Priori?*).

of some sort.¹⁴² In any case though, Sellars maintains the thought that there is *some* connection to be salvaged from Carnap's same-content thesis, even though the connection is not, as Carnap thought, analysis or sameness of content:

Clearly it [i.e. the utterance “ ϕa necessitates ψa ”] conveys (and does not assert) that the speaker conforms to the rule “‘ ψa ’ is L-derivable from ‘ ϕa ’”...

– *Inference and Meaning*, page 281/333.

As it turns out, Sellars thinks that the explicitly syntactic sentence “‘ ψa ’ is derivable from ‘ ϕa ’”¹⁴³ *also* conveys that the speaker conforms to the said rule of inference.¹⁴⁴ In fact, what was in Carnap's hand a tool for *syntaxizing* metaphysical disputes has turned in Sellars hand into a tool for a kind of *psychological behaviorization* of claims about abstract-entities (with the caveat that Sellars does not claim reduction).

Sellars explains the convey vs. assert contrast by reference to the fact that, when a speaker asserts something, say about the weather, his *asserting* act has a significance going beyond *what* he asserts. The act conveys information about the speaker's “state of mind”, i.e. that he believes such and such about the weather. *What* he asserts, on the other hand, is strictly about the weather. “Convey” is reserved for the first sense of an assertion's conveyance of information. While “assert” is reserved for the second sense of asserted content.¹⁴⁵ We can think of this “convey” relation to be what underlies Moore's paradoxical claim. That is, the problem with “It is raining, but I don't believe that it is raining” is that what the second part of the utterance *assert* contradicts what the first part of it *conveys*.

Brandom describes this explanation of Sellars as “only somewhat helpful”¹⁴⁶ and loses no time to give Sellars example a *pragmatist* gloss: “The point, I take it, is

¹⁴² So in *Inference and Meaning*, page 281/333, Sellars writes “In other words, the utterance conveys the existence of a rule-governed mode of behavior in the speaker. But it is equally clear that the utterance “‘ ψa ’ is L-derivable from ‘ ϕa ’”, being a normative utterance, does not *describe* the psychological mechanisms of the speaker” (underlines mine). The flow of the argument makes it clear that Sellars is assuming that “a rule-governed mode of behavior in the speaker” consists of “psychological mechanisms of the speaker.”

¹⁴³ I have dropped “L-” in Sellars' formulation of sentences. “L-derivable” means simply derivable by means of rules of transformation (i.e. of inference) specified as part of the (formal) language L.

¹⁴⁴ The other important instance of this “behaviorization” strategy, is of course the case of meaning statements of the form “E means ...”.

¹⁴⁵ C.f. page 281/333 of *Inference and Meaning*.

¹⁴⁶ C.f. Brandom, Lecture 4 of *John Locke Lectures*.

Part II: Semantic Functionalism & Problems of Laws

to distinguish what follows *semantically* from the content of what I had *said* from what follows *pragmatically* from my *saying* of it".¹⁴⁷ But of course, pragmatic consequence is *not* the point of Sellars' concept of "convey". It is no accident that, both in Sellars' example illustrating the contrast of "convey" vs. "assert", as well as in all the applications that Sellars' makes of the concept of "convey", what is conveyed is said to be *psychological facts*. Psychological facts are not pragmatic consequences, nor are they (the pragmatist concept of) doings. Specifically, facts conveyed by meaning statements, modal statements, as well as statements about norms, are, according to Sellars, facts of the form: the speaker uses the expressions in question in accordance to the rule¹⁴⁸ of inference.... Using expressions in accordance of rules is for Sellars a fact that involves "the conception [on the part of the language user] of the norm enjoined by the rule", and that, is a psychological matter. To quote a passage we have already seen before:

A uniformity in behaviour is rule-governed not *qua* uniformity, ... but *qua* occurring, in a sense by no means easy to define, because of the conception of the norm enjoined by the rule.

– *Inference and Meaning*, page 284/336. (Underlines mine)

What speaks for Sellars exploitation of the concept of "convey" is the fact that one can construct paradoxical statements on analogy with Moore's paradox, at least for the case of normative claims: "' ψ a' is derivable from ' ϕ a', but I do not conform to the rule of inference from ' ϕ a' to ' ψ a'". There is no *logical* contradiction in the content of this assertion: it might be true that *both* ' ψ a' is derivable from ' ϕ a', *and* the speaker does not conform to or reject the corresponding rule of inference, perhaps because he is not aware of the derivability fact. But for the speaker to *say* these two things in one breath *is* problematic.

Brandom's finding Sellars' explanation "dark" is therefore at least partly a reflection, here as elsewhere, of his systematic failure to take seriously the role of the psychological and the mental in Sellars' explanatory framework. It is nonetheless to Brandom's credit that he does not apply the mundane sense of "convey" according to

¹⁴⁷ C.f. Ibid.

¹⁴⁸ Which of course includes certain "conception of norms" on the part of the speaker, hence it is a psychological fact. See §2.4.3 above.

which my assertion conveys my beliefs as *widely* as Sellars applies it. At least the case of *meaning statements* appears to be problematic. For if I say “‘rot’ in German means *red*”, what I thereby convey, according to Sellars, is that *German speakers* use tokens of ‘rot’ in accordance with similar rules as those that govern *my* use of tokens of ‘red’. What is conveyed, in other words, is not just psychological facts about *me*, but a non-psychological relation between various psychological facts (about me and about German speakers). For precisely this reason, we also cannot construct Moore-like paradoxical sentences.

I will leave the question open whether Sellars’ application of the convey vs. assert contrast to *modal* claims is more tenable than his application of it to meaning statements. The reason is that Sellars’ account of modal claims is inadequate on other grounds already (we have already seen indications that Sellars has difficulty with an account of how the concept of modal-connections can ever legitimately be *applied*, in his sense of “apply”), so that the issue of whether a precise sense of “convey” can be found to link modal claims with rule-governed linguistic behavior is ultimately not relevant to our decision to reject the overall account.

Sellars’ concept of “convey” is his answer to Carnap’s same-content thesis. What about Carnap’s non-literalism thesis about necessity statements? Here Sellars follows Carnap more closely, though less explicitly. First of all, in Sellars’ semantic writings there is *general silence* on the topic of what meaning, modal, and normative statements actually *assert*, as opposed to convey. If there is *nothing* that can be said about the asserted content of these statements, then one cannot be blamed for drawing the conclusion that knowing what they convey is *all* there is to know about the significance of these statements. Moreover, there is positive evidence that this is what Sellars actually thinks: conveyance is all there is to be said about the significance of modal statements, among others. For every now and then, Sellars denies in the clearest of words that some of these statements are descriptive of their apparent topics. So in one of the early papers Sellars writes:

But what is *aboutness* but *meaning*? ... To say of an event *e* that it is *about* something is not to describe *e*.

– *Mind, Meaning, and Behavior*, page 92. (underline mine)

Part II: Semantic Functionalism & Problems of Laws

In general, the task of the philosopher is to explore without prejudice the syntactical and pragmatic relationships which obtain between the various forms of discourse, descriptive, semantical, normative, modal, etc. Surely the hankering to give bad marks and pseudo-conceptual status to other forms of discourse merely because they are discerned not to be descriptive discourse belongs with other left-wing deviations in the Madame Tussaud's Way Work Museum of the analytic movement.

– *Ibid.*, page 94. (underlines mine)

In the first of these passages, Sellars claims that statements that attribute semantic aboutness are not descriptive of putative bearers of the aboutness. In the second passage, Sellars widens the contrast to be between descriptive on the one hand, and normative, semantic, and modal, on the other hand. More famously, Sellars denies the descriptiveness of statements attributing knowledge, and I quote again:

The essential point is that in characterizing an episode or a state as that of *knowing*, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says.

– *EPM*, §36. (underline added)

All these passages point to the thought that modal statements, like many statements purported about abstract entities, do not *describe*, they only convey. Their significance can and must be grasped, therefore, by grasping what they convey. What they convey, however, is *not* in turn to be reduced to the contents of modal statements. To that extent, modal statements (as well as semantic and normative statements) are *secondary* to statements about rule-governed behavior. And that, given Sellars' replacement of the same-content thesis by the conveyance thesis, is as close to Carnap's non-literalism about modal statements as one can get.

The semantic upshot of Sellars' view, then, is that first, modal claims do not have *asserted content*, so that, as one is surely entitled to infer, there cannot be a *theory* about their semantic content either. Second, these statements also do not

convey modal facts. They, like normative and semantic statements, convey psychological-behavioral facts.^{149 150}

§3.4 Brandom's Expressivism

A less psychological account of modal statements, one that is still in the same anti-realist spirit of Sellars and Carnap, can be found in Gilbert Ryle and is adopted by Brandom. Ryle's distinction between knowing-that and knowing-how, as well as his idea of "inference tickets" have no doubt also influenced Sellars. But it is in Brandom that we find, for example, Ryle's signature claim about the role of conditionals rise to absolute prominence. Conversely, Brandom's signature concept of explicit-making is an articulation of the *relation* between knowing-that and knowing-how, which is a *further development* of the Rylean distinction.

Before examining the power of the rhetoric of "making explicit" to potentially go beyond Sellars and Carnap, let us first, therefore, turn to Ryle.

§3.4.1 Ryle – "Inference Tickets" & "Hypotheticals"

Ryle's view on the relation between modal statements and inferences is mediated, we might say, by his views on conditionals, which he calls "hypotheticals". So, first, Ryle claims – and Brandom cites him on this with approval – that "the differences between modal and hypothetical statements are in fact purely stylistic".¹⁵¹ In other words, modal claims are stylistic variants of conditionals. Second, Ryle says that conditionals

¹⁴⁹ There is, in effect, some sort of reduction of normative, semantic, and modal facts to psychological-behavioral facts, even though Sellars officially avoids the reductionist charge with his doctrine of "convey". Given the fundamental explanatory status of psychological-behavioral facts for Sellars, it is no accident that he remained a steadfast realist about and developed sophisticated defense for mental items such as thoughts and sense impressions.

¹⁵⁰ To conclude our discussion on Sellars, let me note that the more encompassing version of Sellars' thesis – covering semantic, modal, and normative statements – is potentially problematic. For they are said to convey facts of the form: X uses expression E in accordance to rule R. These facts are therefore stated with the phrase "in accordance to rule R". Sellars would have to, to avoid circularity, maintain that the mentioning of a rule in this phrase is syncategorematic, and that the whole phrase characterizes a psychological feature. But at least some psychological features – those involving *aboutness* – are not literal, for attributions of these do not describe but merely convey. So it seems that Sellars has to say that "in accordance with rule R" describes a *non-intentional* psychological feature. It is neither clear what that would be, nor does Sellars' attempt to make it clear. But these problems notwithstanding, what interests us here is Sellars' account of the *modal* statements, that is, the *narrower* version of Sellars' priority thesis. Consequently, we do not have to deal with the potential difficulty of a non-circular account of intentionality.

¹⁵¹ For Ryle's statement, see his paper "If," "So," and "Because", page 313. For Brandom's approving citation, see Brandom's *John Locke Lectures*, lecture IV, page 17-8.

Part II: Semantic Functionalism & Problems of Laws

are what he calls “inference tickets”, and that they belong to “post-inferential”, rather than to “pre-inferential” levels.

We can illustrate Ryle’s first thesis by saying that the statement “ ϕa necessitates ψa ” is a stylistic variant of the conditional “If something is ϕ , then it is also ψ ”. This is not the controversial part of Ryle’s claim. In fact, the idea that modal statements are all implicitly conditional, and (most) conditionals have a modal force, has become a widely accepted doctrine in linguistics and philosophy of language.¹⁵² The second part of Ryle’s claim goes to the core of our present topic, and can again be divided into two theses, one positive, one negative, corresponding to the two parts of Sellars’ priority thesis:

(R – positive) Knowing “ ϕa necessitates ψa ” or the corresponding conditional is having a license or warrant to make inferences from ϕa to ψa .¹⁵³

(R – negative) Serving as inference tickets is the *only* function conditional and modal statements have. In particular, they do not have a *descriptive* function independent of their instrumental utility for an inferential practice.

The idea is that conditional statements – and by extension modal statements – are just a tool to facilitate the making of inferences. Their only function is an *instrumental* one. In fact, Ryle compares the relation between the making of inference and the assertion of conditional statements to that of *exercising* a skill vs. giving *instructions* about that skill:

In arguing (and following arguments) a person is operating with a technique or method, i.e., he is exercising a skill; but in making or considering a hypothetical statements and explanations, he is, for example, giving or taking instruction in that technique or operation. Roughly and provisionally, he is not cooking, but writing or reading a cookery book, not practising an art but teaching it or receiving tuition in it.

– “If,” “So,” and “Because”, page 331. (underlines mine)

¹⁵² For a locus classicus on this, see Angelika Kratzer’s *What ‘Must’ and ‘Can’ Must and Can Mean*.

¹⁵³ I have adapted Ryle’s formulation to Sellars’ terminology, which is closer to ours. For Ryle’s formulation, see his “If,” “So,” and “Because”, page 308.

That description makes it sound like that the contrast Ryle has in mind, to use another Rylean turn of phrase, is one between “knowing how” and “knowing that”. For, if a conditional/modal statement is like “writing a book” on the subject of the “technique” of making inferences, does not that mean that it *describes that* such and such inferences are good? If that was what Ryle had in mind, then he would be holding the radical Carnapian position that a modal statement has the *same content* as a corresponding statement about syntactic rules.

The fact of the matter is, Ryle is sandwiched between two extremes: (a) the explicit anti-realist stance according to which modal statements have *no* asserted/descriptive content at all, on the one hand, and (b) the Carnapian idealist-reductionism according to which modal statements describe aspects of *our linguistic practice*, rather than aspects of extra-linguistic reality, on the other hand. What brought him to this precarious, dialectically non-stable position, is ultimately his metaphysical scruples. He has the following to say about “necessary connections”, which we saw was for many the “content” of modal statements:

...theorists are apt to ask “What exactly do hypothetical statements assert to characterize what?”... or, more generally, “What do such statements describe?” or “What matters of fact do they report?” And they are apt to toy with verbally accommodating replies about Necessary Connections between Facts, or Internal Relations between Universals, and the like.

– “If,” “So,” and “Because”, page 312.

If conditional, and hence modal, statements do not describe “matters of fact”, and given Ryle’s other, far more plausible, claim that they have an “inference ticket” function, then it is hard to avoid one of the two horns just mentioned: (a) they do not describe *anything*; or, (b) they describe matters of *linguistic practice*.

Ryle’s non-factualism about conditional and modal statements – that they do not describe what they appear to describe – is deeply rooted in his overall philosophical outlook. It is especially connected with his project to show that many kinds of alleged mental occurrences do not exist. In his *opus magnum*, *The Concept of Mind*, Ryle argues for this thesis with a number of strategies, one of which has the following shape. First Ryle argues that to describe a (visible) performance as having a certain intelligent character is not to posit an invisible mental occurrence, parallel to or preceding the observable performance. It is rather to describe the observed

Part II: Semantic Functionalism & Problems of Laws

performance as satisfying a variety of “hypotheticals” that capture the “frame of mind” in which the performance is made. These “hypotheticals” are counterfactual statements of the form “if such and such had happened, the agent would have adjusted her performance thus and so”. In a crucial second step of the argument, Ryle argues that these hypotheticals do not describe “matters of fact”, that their only function is to allow us to make inferences of corresponding kinds. Obviously, for this “mongrel-categorical” strategy against mental occurrences to work for Ryle, it is essential that he subscribes the non-factualism about conditional and modal statements.

Apart from the *methodological* constraint just sketched, does Ryle have argument of *substance* to support his non-factualism about modal statements? The main argument Ryle offers appears to be a version of the Humean argument against necessary connections, and to be based on the assumption – to be rejected in this essay – that inference and observation are the only two ultimate sources of knowledge about matters of fact:

They [i.e. the “addicts of the superstition” that all true indicative sentences “either describe existents or report occurrences”] have to admit, finally, that these postulated processes are themselves, at the best, things the existence of which they themselves infer from the fact that we can predict, explain, and modify the observable actions and reactions of individuals. But if they demand actual ‘rails’ where ordinary inferences are made, they will have to provide some further actual ‘rails’ to justify their own peculiar inference from the legitimacy of ordinary inferences to the ‘rails’ which they postulate to carry them. The postulation of such an endless hierarchy of ‘rails’ could hardly be attractive even to those who are attracted by its first step.

Dispositional statements are neither reports of observed or observable states of affairs, nor yet reports of unobserved or unobservable states of affairs. ...

– *The Concept of Mind*, pages 119-120. (underlines mine)

That Ryle assumes as a matter of course that observation and inference are the only sources of knowledge can be seen by reading the two underlined places in the quoted passage together. The appeal to the unappealing infinite “hierarchy of rails” calls to mind Hume’s circularity argument against using the inference from past regularity to future regularity to buttress the induction method (the inference itself requires induction to justify).

It seems that we will do well to separate Ryle's observation about the function of modal claims as "inference ticket" from his claim that this is the *only* function they have, as well as from the non-factualism about modal claims that underlies the more radical claim. The latter, together with some key presuppositions of the traditional Humean, or "empiricist" arguments, will be debunked by later chapters of this essay. Endorsing the former thesis, however, has no tendency to implicate one in the latter theses.

§3.4.2 Brandom on "Making Explicit"

Part of the reason why we have spent some time looking at the details of Ryle's doctrine on modal statements is that much of it is taken over by Brandom. Brandom cites and approves of, as we saw, Ryle's claim that modal claims and conditionals are "stylistic variants". Brandom, like Ryle, speaks of conditionals as "inference tickets", and agree with him that conditional and hence modal statements belong to "post-inferential" rather than to "pre-inferential" levels.¹⁵⁴ What distinguishes Brandom and Ryle on the topic of modal and conditional statements, is that Brandom makes more of an attempt to say what exactly the *relation* is between modal statements on the one hand, and the inferences they are said to license on the other. Brandom's formal statement is what he calls the "Kant-Sellars thesis", which has the following two parts:

(Brandom – A) In using ordinary empirical vocabulary, one already knows how to do everything one needs to know how to do in order to introduce and deploy *modal* vocabulary;

(Brandom – B) The expressive role characteristic of alethic modal vocabulary is to *make explicit* semantic, conceptual connections and commitments that are already *implicit* in the use of ordinary empirical vocabulary.

– *John Locke Lectures, IV*, page 13-14.

The first part of Brandom's thesis is a clearer, pragmatist paraphrase of Ryle's claim that modal statements belong to the "post-inferential" levels. The second claim goes

¹⁵⁴ For Ryle's statement of this, see "*If*," "*So*," and "*Because*", page 312.

beyond the Rylean metaphor of “inference ticket”. This “expressivist” thesis is the key to Brandom’s thinking about alethic modal statements.

The concept of “making explicit” is central to that part of Brandom’s thought on semantics having to do with *understanding*. But it is also closely related to Brandom’s account of semantic *content*. For this reason, the thesis (Brandom – B) gives us hope that, unlike for Sellars and for Ryle, an account of the *content* of modal statements – as opposed to their intra-linguistic function or indirect implications – might be available to Brandom.

§3.4.2.1 Overview

The notion of “explicit making” has been given many different articulations by Brandom. The most general, but also indeterminate formulation of all is: **to make explicit is to say what one otherwise only does**. The relation of the notion to other aspects of Brandom’s semantic thinking that we have already looked at is this. On the one hand, *to be explicit* is to be in the space of reasons: to be an item that can give reason, and for which reason can be demanded. Brandom’s explanation of linguistic *understanding* describes it as a matter of having the capacity for, and participating in, this “game of giving and asking” for reasons, especially in the inferential form. So the notion of *being* explicit is bound up with Brandom’s theory of understanding.

On the other hand, *what* is made explicit – in the sense of the *result* of explicit making, according to Brandom, is a content. We shall yet see whether the explicit-making relation itself is *constitutive* of the content made explicit. But *if* so, we will have a new form of semantic explanation that does not appeal only to inferences, strictly so-called. For the explicit-making relation cannot be a variety of inferential relation, for the simple reason that, one of its relata – the non-explicit doing or some aspect of it – is not even a claim or statement, or anything that stands in the space of reasons. It is this hope of a broadened form of semantic explanation that I shall exploit in the rest of this essay.

Turning now to Brandom’s discussions of the idea, consider first the way Brandom explains the idea for the first time:

Explicitation is capturing as the content of a claim a discrimination of the appropriateness of various performances which otherwise is expressible only in practice.

– *Varieties of Understanding*, page 42.

This needs a lot of unpacking. First of all, making explicit is a *logical* relation, rather than, say, a causal relation, or a kind of spatial-temporally locatable actions, such as the drawing of an inference is. One relatum of this relation is a kind of doing, but not *merely* a doing: it is also an act of *asserting*. This is the *explicit* side of the two relata. A claiming act is explicit in virtue of the fact that such an act can serve to give reason and for which reasons can be demanded, or simply, it is right “in the space of reasons”.

There are difficulties both about the *other* relatum of this relation – that which is *made* explicit, call it the *explicitandum* – as well as about the explanatory role of this relation itself. The relation of explicit making obtains when one of its relata, i.e., a claiming act (or a type of claiming acts), “captures as its content” the other relatum, i.e. the explicitandum. The main difficulty concerning the relation of explicit making is whether and how it can figure in an account of *content* of the explicating act. The “capture as content” formulation can give the impression that an assertion stands in an explicit-making relation to something else *in virtue of* the content it has independent of that explicit-making relation. But as I shall argue, this is an uninteresting way of understanding the idea. The main difficulty with the explicitandum is simply to determine *what* kind of thing it is. Let us give names to these difficulties for future reference:

(Explicitandum Question): *What kind of thing is made explicit? What, in other words, is “captured as content”?*

(Content Question): *Does the explicit-making function of a claim enter into an account of the content of that claim? If so, how?*

On the face of it, Brandom has provided answers to the Explicitandum Question: *aspects of doings* are made explicit. But which aspects, and what kind of doings? It is here that Brandom’s answer begins to vary from context to context, and even within the same context. Let me now continue my discussion by taking a close look at the two examples of explicit-making: observation reports and logical vocabulary.

§3.4.2.2 *Observation Reports*

Brandom's saying-versus-doing formulation takes on a bit more specificity in the case of observation reports:

The responsive performances of both the instrument and the observer have a significance for others But the observer does and the instrument does not grasp or attribute such a significance to its own response. It is in virtue of the explicitness *to* the one who produced them *of* the significance of those responses that the observer's performances count not only as doing something but as saying something. Explicitation of the first kind to be discussed consists in this coming to be able to express in *saying* the practical capacity which otherwise remained an implicit feature of one's *doings*.

– Ibid., pp.32-3, (Underlines mine)

Part of what Brandom is doing here in this passage is to elaborate on the notion of “explicitness”. A performance is explicit only if its significance is *grasped* by the performer. Grasping the significance of a performance, in turn, is according to Brandom being capable of navigating around it in an inferential space. What makes an observer's observation report *explicit*, or, a matter of *saying*, is that the observer is capable of drawing inferences from that report. But what is in the case at hand *implicit*? Which feature of which doings is *made* explicit?

It is more or less clear that the doings in question are “responsive performances”: the conditioned and reliable responses to the environment that classifies it (according to color, or temperature, for example). Bracketing the question about the implicit feature of these doings for the moment, several alternative interpretations open up already before us. According to the first interpretation, the doings whose certain aspects are made explicit by observation reports are similar reliable responses (“similar” in point of issuing from reliable differential capacities¹⁵⁵) by *instruments*, or otherwise non-rational agents. We might say that the explicit-making in this case is *heteronomous*. According to the second, *autonomous* interpretation, the doings in question are precisely the very *same* verbal performances

¹⁵⁵ Brandom's phrase “reliable differential responsive *dispositions*” is a careless confusion between *abilities* and *tendencies*. Dispositions are tendencies, and, like the latter, have a *must*-type of modality. Capacities and abilities have a *can*-type of modality. What is required for a speaker to make observation reports is not that he is *disposed* to uttering “red” when looking at red things, but that he is *able* to do that.

that are the observation reports, but *qua* performances issuing from reliable responsive dispositions. An observation report *qua* a *saying* makes explicit – autonomously – an aspect of the very same verbal performance *qua* a manifestation of a reliable responsive disposition.

There are other interpretations. I will mention one that will become important for logical vocabulary. This is what I shall call the *developmental* interpretation. It has in turn an ontogenetic and a phylogenetic variant. According to this interpretation, the capacity to make observation reports can be reconstructed in terms of two stages. In the first stage, one learns to make reliable responses that classify the environment in some fashion. In a second stage, one learns to make inferences that connect the kind of verbal responses that one has learned to make in the first stage. Diachronically speaking, the same phonological type of verbal response morphs from a mere reliably classifying response, a *doing*, into a *saying*. So, a school boy's utterance of "that's red" makes explicit in this developmental sense certain aspect of earlier utterances of the same type by him when he could not yet draw any inference from it.

Brandom's exposition is never precise enough for us to determine which interpretation he intends. But all three interpretations we have canvassed have something in common: the *capacity* for the doings in question is a component of the capacity required for the very saying that makes these doings explicit, even though the *doings* themselves need not be, depending on interpretation, identical with the sayings. For, according to any of these interpretations, one cannot make an observation report without having reliable responsive capacities¹⁵⁶ for the relevant types of classificatory verbal responses, though having such capacities is not enough to count as capable of making observation reports. This is an important feature of the explicitandum because, as we shall see presently, that it opens up the possibility of an account of the contents of observation reports in terms of their explicit-making role.

In any case then, the doings in question are *reliable responses that classify the environment*, or, to make use of Brandom's terminology, reliable differential responses.

¹⁵⁶ C.f. footnote 155 for the distinction between "capacities" and "dispositions".

§3.4.2.2a “Discrimination of Appropriateness” Made Explicit

But which *feature or aspect* of the conditioned reliable differential responses is made explicit by observational reports? Some of Brandom’s descriptions are clearly problematic. In the passage just quoted, he speaks of “saying practical capacities”, which phrase, understood literally, is just absurd. Brandom’s first general characterization of explicit-making that we quoted above talks about “discrimination of appropriateness”. One problem with this is that it is hard to see how reliable differential responses, found both in human reporters as well as in measuring instruments, can have “discrimination of *appropriateness*” as an aspect. But suppose we allow that Brandom can talk this way because he has a specific picture in mind, say, the *autonomous* interpretation mentioned earlier, according to which the eligible doings to be made explicit are not just *any* reliable differential responses, but those very same responses that qualify as observation reports. And let us assume that a *reporter’s* reliable differential verbal response *can* be thought of as discriminating a particular verbal response as appropriate.¹⁵⁷ We are still left with an unwelcome result: what is made explicit and “captured as content” in an observation report is *the reporter’s discrimination of appropriateness!* For one thing, it is not immediately clear what this means. Is the explicitandum an *appropriateness*, so that the content of the observation report is deontological? Or is the explicitandum the *act* of discriminating, so that the explicit content is about human behavior? Or perhaps the explicitandum is the reporter’s *capacity* for such acts, so that the explicit content is about a human capacity? The more serious problem is that this formulation leads, no matter which of these ways we choose to understand it in, to a sort of non-factualism according to which observation predicates like ‘red’ are not about extra-linguistic reality, but about “oughts”, human behavior, or human capacities. If this is what explicit-making brings us for observation predicates, prospects for a realist semantics for law statements with the help of explicit-making seem pretty dim.

§3.4.2.2b “Information” Made Explicit

A much more promising place to reconstruct a position is where Brandom says that it is the *content* that is made explicit:

¹⁵⁷ That is not an uncontroversial idea. We say that the reporter discriminates a *feature in the environment*, not the appropriateness of his words. But I will let it pass for the sake of the argument.

§3 Inferentialists On Nomological Necessities & Explicit Making

(Explicitandum Thesis) The repeatable content is what is made explicit.

– Ibid., page 34.

I want to explore, not the *resultive* reading (the *result* of explicit making is a content), but the *relatum* reading (the explicitandum of explicit-making is a content). First note that under the *relatum* reading, this is precisely what Brandom officially *cannot* say. The reason is that having a content, according to Brandom, *cannot* be a feature of mere reliable classificatory responses *as such*. Having a content entails being in the inferential space, and so is a feature of the *explicitating* saying, not a feature of the doing *explicitated*. Yet the line quoted cannot be thought of as an isolated carelessness by Brandom. For Brandom frequently uses related terms in explaining the explicitandum, including “significance” and “information”. One gets the impression that “information” in these contexts is the idea of an *inchoate* form of content, so that we might modify Brandom’s claim to:

(Explicitandum Thesis): The repeatable proto-content is what is made explicit.

with some yet to be specified notion of *proto-content*. Here is a telling passage in which some of related descriptions such as “significance” and “information” appear:

...the content understood codifies explicitly (as the content of that claim) the merely discriminative classificatory significances implicit in the corresponding non-assertional differential responses. For information is expressed in the differential responsive behavior of a measuring instrument in a “language” (to stretch that word beyond its proper inferential assertional application) defined by the inclusion and exclusion relations between its response kinds.

– Ibid., page 36. (Underlines mine)

This passage also gives some clue to how we should think of the *proto-content*, or “information”. The idea is presumably this. Each reading of the measuring instrument corresponds to an observation report. There are *factual* entailment- and exclusion-relations amongst various possible readings of the instrument, which corresponds to inferential relations amongst corresponding observation reports. These factually existing relations “define” an “information” carried by each reading. Since the measuring instrument is incapable of drawing inferences, this information as well as the factual relations that define it remain “implicit” to the measuring instrument. For

the observational reports, however, the relations that impart significance are not merely *factual* but *inferential*, as people who make these reports are capable of drawing inferences from and to these reports, corresponding to the factual relations connecting the readings of the instrument. For this reason, a reporter's words differ from instrument readings in that their *contents* are "explicit" to her. But *modulo this difference*, the "information" carried by a reading is *identical* with the content carried by a corresponding report.

There are two "technical" points to note from this passage about the notion of "information", or as I prefer to say, of *proto-content*, before we move on. First, with "information", Brandom does not mean *singular* propositional (proto-)contents that vary from one instance of the report "that's red" to another. Instead, the "information" is *repeated* in each instance of the report "that's red". For the inclusion and exclusion relations that define such information obtain between various reading *types*, or mere response *types*. This also explains the qualification "repeatable" in Brandom's formulation of the Explicitandum Thesis.

The second point is that, for Brandom, the individuation of proto-contents carried by instrument readings, or mere reliable differential responsive capacities, *parallels* the individuation of corresponding content: the former by *factual* inclusion/exclusion relations, the latter by structurally parallel *inferential* relations. Moreover, both are *holistically individuated*.

§3.4.2.3 Revolutionary Rhetoric

What we see here is an implicit tension between two formulations of explicit-making for observation reports. The formulation that takes explicitandum to be "discriminations of appropriateness" leads us straight away back to the kind of non-factualism of Sellars and Carnap. Though at first for observation reports, non-factualism for nomological statements results, as we shall see, in exactly the same way when we stick to this formulation of explicitandum for conditionals. But this is the side of tradition that Brandom inherits from Sellars and Ryle. The formulation of explicitandum in terms of "information" manifests another side of Brandom's story. This is the side of the independent, intrinsic power of the rhetoric of *making explicit*.

The tradition Brandom inherits does not credit the *doings* – "discriminations of appropriateness" for example – with any sort of content. These are *merely* knowing-how, not knowing-that. Only a knowing-that has a content. But the rhetoric

of *making explicit* carries with it the naturalness of talking about *contents implicit in these doings* that are then made explicit in sayings. Brandom succumbs to, or exploits, though it is not clear which, this aspect of the rhetoric, when he says that an observation report makes explicit the “information” carried in mere reliable differential responses.

Since the tradition side of Brandom has only non-factualism to offer for nomological statements, it is the agitating power of the rhetoric of *making explicit* that appeals to the naïve realist. For, with the “information”-formulation of explicitandum, we are finally able to conceive a *route* for empirical information to reach the subject: first the information comes to the subject *unconsciously*, in the form of an “information” encapsulated in his *doings*; then in an expressivist step, the subject *makes explicit* this information in the form of a claim. This is putting things in *epistemological* terms. For nomological statements, however, this epistemological route is acceptable only to someone who already accepts the intelligibility of nomological information.¹⁵⁸ In Part II we shall see that the way to exploit this epistemological route is to describe a coherent practice in accordance with it *without* appealing to the notion of nomological information. Once we secured the presumption of contentfulness for nomological statements, we can *then* describe a step in that practice as doings carrying nomological information.

Epistemology and coherence of nomological practice are topics of later chapters. In semantic terms, and for observation reports, we can examine the implication the “information”/proto-content formulation of the explicitandum by examining the options it allows for the Content Question. These are as follows. *Either* the fact that ‘red’ means what it means is explained (partly) by reference to the fact that it makes explicit a *proto-content* (information) carried by a corresponding instrument reading. *Or*, the content of ‘red’ is explained only by reference to the inferential relations that govern its use, in which case the fact that ‘red’ *happens* to make explicit a certain *proto-content* plays no role in explaining why ‘red’ means *red*. It appears that Brandom has no intention to appeal to *proto-content* to explain the meaning of ‘red’. In fact, in later writings, starting from *Making It Explicit*, Brandom drops the talk about “information” carried by mere reliable differential responses altogether. In a

¹⁵⁸ In fact, the very concept of *information* would seem to presuppose nomological regularities in the world.

Part II: Semantic Functionalism & Problems of Laws

way, Brandom also has no reason to appeal to proto-content, given that the way the constitution of proto-contents is functionally explained parallels *exactly* the functional explanation of explicit contents. Instead of a (non-normative) functionalist explanation of proto-content *plus* a theory of explicit-making, Brandom can, and does, directly tell a (normative) functionalist story about explicit contents.

I think that the *horizontal* kind of content explanation, in terms of holistic inferential or inclusion/exclusion relations, is ultimately incapable of accounting for the *empirical* nature of the content of observation predicates such as ‘red’. For this reason, I believe the *other* answer to the Content Question is more promising, according to which the explicit-making relation is partly constitutive of the meaning of ‘red’. Let me call this a *vertical* kind of content explanation. But if a vertical explanation of content is going to avoid holism, the *proto*-content made explicit by ‘red’ must not be thought of *horizontally*, as defined by inclusion/exclusion relations of mere reliable responses. A promising theory will take *proto*-contents to be carried by *perceptual experiences*, and perceptual judgments to be acts of explicit-making of such proto-contents.

I cannot go into detailed argument against holistic theories of content/meaning, nor detailed argument for a vertical theory of content for observation predicates,¹⁵⁹ each of which requires book-length treatments. A question more relevant for us, is: is a vertical content explanation that makes the explicit-making relation account for the content of the explicit assertion, possible for nomological modal statements? As I indicated earlier, and shall argue below, such a semantic explanation, referring to nomological information as it does, can only come *after* the presumption of contentfulness for nomological is secured. And the latter, requires dealing with epistemological questions about laws.

The most revolutionary aspect of the *explicit making* rhetoric, then, is its suggestion of an epistemological path to laws. Following this path requires a delicate dialectic strategy, and another idea (that of harmony). But we should hold on to the idea that doings (or experiences) can carry implicit *empirical* information. For that

¹⁵⁹ One important modification of the idea of explicit-making needed for such a theory is that explicit-making is not just a matter of transposing *exactly the same* information from one medium into another. A genuinely *expressive* notion is needed. The two main advantages of such a theory are, (a) such a theory avoids the perils of content holism, and (b) the role of *experience* in such an account also helps Brandom to meet challenges to his account of *understanding*, such as those leveled by McDowell we saw in §2.4.4.2 above.

idea also – under adequate metaphysical assumptions – makes a new kind of functionalist semantic explanation possible, one that relies not only on intra-linguistic use, *inferences*.

Let me now turn to Brandom's expressivist account of logical vocabulary. Some of the divisions and tensions are the same there. What I want to focus on is whether the expressivist strategy can perform some sort of semantic bootstrapping: making sense of nomological content *without* an antecedent resolution of the traditional, epistemological problems about laws.

§3.4.2.4 *Logical Vocabulary*

If the explicit-making aspect of observation predicates is the aspect of *understanding*, the explicit-making aspect of logical vocabulary is the aspect of *self-understanding*. The promise to give an account of self-consciousness is precisely the charm about the notion of explicit-making that makes it an enduring theme of Brandom's writings, right up to the present time.

The interest of the putative expressivist function of logical vocabulary for *us*, however, lies not in a theory of (self-)understanding, but in the possibility of a realist theory of modal *content*. Logical vocabulary is especially relevant because, as we saw, Brandom agrees with Ryle that modal vocabulary is a "stylistic variation" of the conditional, and the conditional happens to be Brandom's favorite example for the expressivist function of logical vocabulary. We have seen that Sellars and Ryle have difficulties acknowledging that modal claims have *empirical* contents (they swing between the no-content view to the view that the content is about an aspect of our *linguistic* practice). So our question is this: does Brandom's expressivism allow us to maintain the obvious, that modal statements describe aspects of extra-linguistic reality?

Let us start with the first of the two questions we posed about expressivism earlier. Which aspect of which doing is made explicit by the conditional statement "if p, then q"? According Brandom, the doing in question is the *practical endorsement of the inference, or of the inferential propriety*, from p to q. But which aspect of that is "captured as content" in the conditional statement? For reasons similar to those we found in the case of observation reports, it is implausible to say that it is an *endorsing act* – that is, an act of inference making – or the *inference pattern* endorsed, or perhaps the (*inferential*) *propriety* endorsed, that is "captured as content". For then we

Part II: Semantic Functionalism & Problems of Laws

would be saying that a conditional statement, and consequently also a modal statement, has a content that is deontological, or about inference patterns or inference-making activities. The case here is not as clear-cut as the case of observation reports, for conditional and modal statements are extremely heterogeneous, and some of them are indeed not plausibly regarded as being about features of extra-linguistic reality. But the problem is that there *are* nomological modal statements, and these would be treated – under such an understanding of “explicit making” – as having contents about various aspects of our *language*, rather than about the world. If so, we are not further along than Ryle, Sellars, or even Carnap.

§3.4.2.4a The “Nomological Conditionals”

Let me explore for a moment the idea that the kind of conditionals of which we intuitively think as describing nomological aspects of reality – the nomological conditionals – make explicit aspects of inference-making practice. The way to move forward is to understand the doings here, namely, acts of endorsing inferences of a certain type, as *carrying a proto-content* and to see this proto-content as what is made explicit in a conditional claim. My treating the inference from “there’re dark clouds moving in” to “it’s going to rain” as good carries something that is *content-like* because one can extract (correct or incorrect) information from this endorsement of mine. Yet what it carries is only a *proto-content* because this endorsement itself does not have the form of an *assertion*, it is not something whose manifestation can count as reason giving, nor can reason be directly demanded for such an endorsing act. (If I say: “there’re dark clouds moving in; so, it’s going to rain” you cannot question the goodness of my *inference* by asking “why do you say that?” or “What makes you think that?”. To challenge, you would have to say “why do you think that *if* there’re dark clouds moving in, *then* it’s going to rain?”.)

In fact, Brandom can be thought of having a picture of this kind where the *proto-content* is the information that a certain form of inference is appropriate. We want to see if this picture can be had with a narrow class of conditionals, and with *proto-contents* that are *empirical* (nomological connections), rather than deontological.

Given this picture, what about the Content Question? How does the explicit-making function of a conditional or modal statement play a role in constituting its

content? Here the situation differs from Brandom's official account of observation reports. The treating of certain inferences as good is numerically distinct from the making of the corresponding conditional statement, whereas a report *is* an exercise of the reliable differential responsive capacity whose proto-content is made explicit by that very report. In other words, the way explicit-making for observation reports plays a role in semantic explanation is through the *numerical identity between the doing and the saying*. For conditional statements, however, *the implicit endorsement of inference is the condition of application as well as the consequence of application for the explicit asserting of a conditional*. One way of describing the difference is this. What a reliable differential response requires for the proto-content it carries to become content proper is just that it, the very response, be imbedded in an inferential practice (among other things). By contrast, the proto-content carried by my endorsement of an inference pattern has to be *extracted*, as it were, by another, numerically distinct act that is a statement. What we want to know is whether the following *positive answer to the Content Question* is possible: the explicit-making relation is constitutive of the content of the conditional sentence.

Now endorsements of inference patterns can, but need not to, carry proto-contents that have an empirical origin. Those made explicit by *nomological* modal claims, however, must carry empirical proto-contents if we are to deem the former as describing aspects of extra-linguistic reality. Let me call such endorsements the endorsements of *nomological inferences*, on analogy with “nomological conditionals/modals”. Nomological inferences are then a species of material inferences. To give a complete account of how nomological conditionals acquire the contents they have, we need to give an account of how endorsements of nomological *inferences* acquire the *proto-content* they have. It is here that a promising expressivist strategy encounters a problem. Let me explain.

§3.4.2.4b Do Nomological Inferences Have Proto-Content?

To explain how the endorsements of nomological *inferences* have the nomological *proto-contents* that they would have to have requires, it appears, an account of how these inferences and their endorsements are connected with some *interactions with nature*, assuming that the nomological *proto-contents* have empirical origins. This interaction need not be *experience* in the classical sense, but *some* form of interaction

Part II: Semantic Functionalism & Problems of Laws

enabling the information transfer must exist. (For comparison, Brandom denies that the “information” implicit in the reliable responsive dispositions to classify is to be explained in experiential terms. But even there we have to assume that the reliable responsive *disposition* is developed by interacting with nature. Suppose, first, that nomological inferences get their *proto*-content by being *justified* by something else. That would defeat the purpose of accounting for nomological content via the *proto*-content of inferences. For the supposed justifier of such inferences, which expresses a proposition, must already have nomological content. And if so, we might as well directly ask how *statements*, such as these justifiers, come to acquire nomological content, forgetting the question how certain *inferences* come to acquire nomological *proto*-content. Moreover, endorsements of inferences are not even the right *kind* of thing for which evidence can be provided (by contrast, claims have evidence). Suppose then the endorsement of nomological inferences acquire *proto*-content in virtue of an information transmitting causal relation to an interaction with nature. It is hard to imagine what sort of interaction with nature can reliably *cause* me to endorse an inference pattern in a way that *information* is thereby transmitted. The nomological *connection* itself is not the right sort of thing to *cause* anything. For it is not a *particular* (event), so cannot stand in causal relation to anything. Even if we let this problem pass, how can a *single event* transmit, by causing some linguistic behavior of mine, the *general* information of a nomological regularity?¹⁶⁰

The problems appear many and serious. But the *real* problem is that the discussion of the expressivist strategy so far puts the cart before the horse: it takes over Brandom’s assumption that nomological conditionals/modals make explicit (endorsements of) nomological inferences. That is not the correct way to account for nomological contents. Quite the contrary, nomological inferences do not provide warrant for nomological statements: they are justified by them. Expressivism can indeed be applied to nomological statements. Yet *what* they make explicit is not, as I shall argue in §5 and §7 below, any aspect of the inferential practice, but the *proto*-content in certain kinds of practical *skills*. But even this correction of the doings that actually bear nomological *proto*-information does not promise a direct semantic bootstrapping. Such a semantic explanation still makes reference to nomological

¹⁶⁰ Mismatch between singularity of evidence and generality of claim is a core aspect of Hume’s anti-realist arguments. To think of singular causation as capable of transmitting general information will encounter similar difficulties.

information, and will be felt as circular for someone who is suspicious of the very reality of laws.

The delicate dialectic movements required to cure skeptical urges are made in Part III below. There I shall tackle what I take to be the most fundamental obstacle for seeing to it that the realist intuition about laws gets the philosophical approval it deserves. I shall return to explicit-making as a mode of *semantic explanation* in §7.

Appendix to §3.4.2.4

To end the discussion, let me go back to Brandom's expressivism about logical constants. What we have just seen is that at least for *one* kind of logical constant – the nomological conditional forming device, insofar as such a thing can be identified – the expressivist strategy of Brandom does not work if we want to preserve our realist intuitions about laws. It is however arguably true that for a *general* kind of conditional, the conditional forming device makes explicit the subject's endorsement of a corresponding inference. Here "explicit making" has a strong Sellarsian taste, à la his notion of "convey".

For other logical constants, it is not clear how far Brandom's expressivism can go. The identity sign, for example, is said to make explicit the speaker's endorsement of certain substitutional inference patterns. But this is a very counterintuitive claim. For it seems that the subject's endorsement of a substitutional inference should be explained by his endorsement of the corresponding identity statement, not the other way around. Another kind of problem is presented by constants like the conjunction sign. What kind inferential propriety, or endorsements thereof, is expressed by conjunction? In the *John Lock Lectures*, Brandom seems to adopt of the analysis of conjunction in terms of other logical constants. But it seems not very plausible that simple conjunction makes explicit – in a way that is somehow constitutive of self-consciousness – some complex function of proprieties associated with other logical constants.

To evaluate Brandom's expressivism for logical vocabulary in detail would require another book-length essay altogether. I will conclude now by drawing the preliminary conclusion that it is not clear whether Brandom's expressivism about logical vocabulary can be successfully carried out.

§4 Harmony – A Diagnosis

§4.1 The Idea Of Harmony	155
§4.1.1 Harmony as a Necessary Condition for Meaning Fixation	155
§4.1.2 Logical Constants vs. Non-Logical Vocabulary	157
§4.1.3 Aspects of Use Essential to Meaning.....	158
Appendix to §4.1.3	160
§4.1.4 Total Harmony Or Intrinsic Harmony?.....	161
§4.1.5 Necessary vs. Sufficient Condition.....	163
§4.1.6 Presumption of Unique Meaning.....	164
§4.2 The Problem of Laws of Nature	165
§4.2.1 Overview.....	165
§4.2.2 Intensionality of the Law Operator and Two Types of Accounts.....	167
§4.2.3 Apparent Disharmony I – The Problem of Induction	169
§4.2.4 Apparent Disharmony II – The Problem of Counterfactuals.....	172
§4.3 Realism vs. Anti-Realism About Laws of Nature	174
§4.3.1 The Anti-Realists	174
§4.3.1.1 Hume	174
§4.3.1.2 Contemporary Anti-Realists	175
§4.3.2 The Realists	178
§4.3.2.1 The General Position	178
§4.3.2.2 The Criterial View – Transition To A Use-Theoretic Account	181
§4.3.2.2.1 Defeasibility and Ontological Gap.....	182
§4.3.2.3 Realism Backed by Analysis and Truth-Conditions	185
§4.3.2.3.1 Necessitarian Analysis.....	185
§4.3.2.3.2 Universal Analysis	188
§4.3.2.3.3 Why Realism-by-Analysis Fails.....	190

§4.1 The Idea Of Harmony

There are two specific ideas I borrow from the inferentialists in my current pursuit of an adequate semantic account for nomological claims. Explicit making is one. The other is the idea of *harmony*. It serves as a crucial instrument for *diagnosing* the root of the perennial difficulty modern philosophers have had with nomological statements and facts. The goal of this section is to introduce and critically examine this notion as Dummett employed it, and adapt it for use in the diagnostic work to be done in the rest of the chapter.

Harmony is a *condition of linguistic use for conceptual coherence*. What this means is that it specifies a condition on the way we use an expression for that expression's being coherently expressive of a concept. This formulation is imprecise and vague as it stands. One of the issues is that not all aspects of an expression's use are essential to its *meaning*, so some must be excluded from the requirement of harmony. Another issue is just how strong we want the requirement of harmony to be: a stronger notion of harmony requires – whereas a weaker variant does not – *uniqueness* of the concept expressed given the rules of use. Let me discuss these details by reviewing the notion of harmony as Dummett and Belnap first introduced the concept, for logical constants.

§4.1.1 Harmony as a Necessary Condition for Meaning Fixation

The historical background to Dummett's notion of harmony is an essay written by Prior and a response by Belnap on the feasibility of conceiving logical constants as having their meaning completely fixed by deduction rules in which they appear. Prior constructed an artificial “logical connective”, called “tonk”, with an introduction rule identical with that of the disjunction sign and an elimination rule identical with that of the conjunction sign. The addition of “tonk” leads immediately to the provability of contradictions, which shows that the said combination of use-rules does not determine a meaningful connective. Belnap responds to Prior's challenge by stating that the introduction of a connective by deduction rules must fulfill two conditions: (i) the addition of the new connective and the rules results in a *conservative* extension, (ii) these rules *uniquely* determine the concept expressed, in the sense that if a second connective were introduced with the same set of rules, then both connectives added

Part II: Semantic Functionalism & The Problem of Laws

would be provably equivalent with each other. The notion of a conservative extension of a language is this: the addition of some expressions and a number of associated use-rules to a language is *conservative*, just when no deductive inference with premises and conclusion exclusively formulated with the old, unextended language becomes valid that was not valid before the addition. In short, a conservative extension does not allow one to prove more things in the old language than one could already.

Dummett picked up this exchange. In fact, “conservativeness of extension” will be one of his formulations of the requirement of harmony. But Dummett’s formulations and discussions are done in the wider and more general context of his use-theoretic semantic theorizing, and for that reason, will be more of interest to us. Dummett distinguishes, quite generally for any natural language, two aspects of linguistic use. On the one hand, there are *principles of verification*, which are meant to include both principles about “when an assertion is conclusively established” as well as those about “what (merely) warrants the assertion’s being made, though defeasibly”.¹⁶¹ As examples Dummett mentions “reports of observation” that are arrived at without inference but based on perceptual experience, as well as mathematical theorems established by proofs. On the other hand, there are *pragmatic principles*, which have to do with the *consequences* the acceptance of a statement has for a speaker, which may either be another statement, or an action. Dummett says that the notion of a harmony requirement is rooted in the intuitive but compelling idea that these “two complementary features of any [language] ought to be *in harmony* with each other”.

Note that the distinction of the two aspects of linguistic use by itself does not force one to accept *verificationism about meaning*: one endorses that only by committing to the additional claim that the meaning of an expression is constituted *entirely* by the verification side of its use. In fact, as we shall see later, the distinction need not, *by itself*, even commit one to *use-theoretic* semantics, for it does not identify meaning with the two aspects of use so distinguished.

Dummett, however, uses the distinction to formulate – *assuming* that the meaning of a logical constant or indeed of any expression whatsoever is fixed entirely

¹⁶¹ C.f. Dummett’s *The Logical Basis of Metaphysics*, chapter 9, page 211. The following quotations are from the same chapter unless noted otherwise.

by ways of use deemed appropriate by competent speakers – a condition of success for fixing the meaning of a logical constant by stipulating deduction rules. This means that Dummett is ultimately interested in a *sufficient* condition for meaning determination. Harmony may not be the whole answer, but it is part of the answer.

This is not how I shall exploit the idea of harmony in this essay. Harmony will serve only as a diagnostic tool, and as such, will only take the form of a *necessary* condition of conceptual coherence, rather than a sufficient condition for meaningfulness.

The easiest way to describe the idea of harmony as a necessary condition for meaning fixation, is to say when it is *violated*:

If there is disharmony [between the two aspects of use involving an expression **E**], it must manifest itself in consequences not themselves involving the expression **E** but taken by us to follow from the acceptance of a statement **S** containing **E**. Acceptance of **S** must issue directly in actions not warranted by the grounds, rated adequate under our linguistic conventions, on which the statement **S** had been made; or our having accepted **S** as true might be taken to justify some further assertion not involving **E**, likewise not warranted by the grounds for making **S**.

– *The Logical Basis of Metaphysics*, page 218 (underlines mine)

To put it simply, there is disharmony between two aspects of linguistic use when the pragmatic principles allow us to draw from a statement consequences, in words or in action, not warranted by what the verification principles deem as adequate ground for that statement. Though Dummett himself makes systematic application of this idea only in the case of logical vocabulary, the generality of the above formulation will be essential for us later. Let me now examine various dimensions in which the principle of harmony can be made more precise.

§4.1.2 Logical Constants vs. Non-Logical Vocabulary

When limited to logical constants, the distinction of two aspects becomes easier to grasp. In that case, both sides are made of *intra-linguistic* principles. The verification side consists of principles linking a statement containing a given logical constant with

Part II: Semantic Functionalism & The Problem of Laws

other statements that deductively establish the former. The pragmatic side consists of principles linking that statement with statements that follow deductively from it. These are better known as, respectively, introduction and elimination rules for logical constants.

It is the case of non-logical vocabulary that will eventually be of interest for us. For non-logical expressions, each of the two aspects of use will be mixed. So the verification side of use for ‘red’ will include *not only* intra-linguistic principles like: “X is red” can be conclusively established by “X is crimson”. It can, and probably should, *also* include such principles as: The assertion “X is red” is defeasibly warranted by a visual experience of a certain type on the part of the speaker. The latter principle is not an intra-linguistic one, because it cannot be reformulated as: “X is red” can be defeasibly established by “I have had/am now having an experience of the kind *red*”.¹⁶²

The pragmatic side of use for non-logical vocabulary can be a bit complicated. Like the verification side, it also has a mixed constitution. On the one hand, the pragmatic side of use for “red” includes *intra-linguistic* principles such as: accepting “X is red” obliges one to accept “X is warm colored”. It also includes *non-linguistic* principles concerning actions such as: given that one wants a red scarf, accepting “X is a red scarf” justifies, defeasibly, the action of trying to get X. The complication here lies both in the specification of an existing *desire* or some other conative state, as well as in the nature of “justification” involved.

In my application of the harmony principle to nomological statements, I will have a lot more to say both about these complications on the pragmatic side of use, as well as about the nature of *experience* that figures in the verification side of use.

§4.1.3 Aspects of Use Essential to Meaning

The next qualification to be made is that the extent of the two aspects of use subject to the harmony principle should be limited to that which is *relevant to meaning*. The idea can be illustrated by the conjunction sign “and”. It is arguably a part of our use of

¹⁶² For one need not to *reflect* on the fact that one is having the experience of a certain kind to have the warrant based on experience. Moreover, the principle of warrant just cited is just a *form* of such a possible principle, and as such, has an *externalist* formulation. The form of the verificationist principle requires only the *fact* of perceptual experience, but does not specify their *nature* nor the extent to which they are conceptualized. But particular versions of the principle need not preclude additions of internalist requirements.

that expression that, when we hear “A called and he ate”, we feel entitled to draw the conclusion that A *first* called, and *then* ate. The “verification principles” involving “and”, on the other hand, allow us to establish the conjunction by establishing each of the conjuncts separately, without paying attention to the temporal sequence of the eating and calling they each describe. Obviously, if this description of use is correct, such verification-grounds do not entitle us to draw the conclusion we feel entitled to draw about temporal sequence. Consequently, the two aspects of use involving “and” are out of harmony. One response available to us is to say that the “consequence” of temporal sequence is not a part of use *relevant to meaning*, strictly so-called, and that it pertains instead to a possible *implicature* of conjunctions.

Another example is afforded by a non-logical expression discussed by Dummett briefly elsewhere. The expression in question is the pejorative term “Boche”, where “the condition for applying the term to someone is that he is of German nationality; the consequences of its application are that he is barbarous and more prone to cruelty than other Europeans.”¹⁶³ The disharmony between these two aspects of use consists in the fact that the consequence of the statement “X is a Boche” sanctioned by the pragmatic side of use goes beyond what is warranted by what the verification side of use specifies as adequate ground for asserting it, namely, the statement “X is of German nationality”. Now one possible analysis of this case is that the Alsatians who used the term “Boche” around the time of World War II had a negative *emotional association* with Germans that is expressed by the pragmatic side of the use they gave that term. So the fact that the conclusion of cruelty and barbarity can be drawn from being a “Boche” is not relevant to the *meaning* of that term, but rather is expressive of a psychological association on the part of the person who uses it.

The point of these examples is not, say, about how properly to analyze pejorative expressions.¹⁶⁴ The point is rather that it is quite conceivable that not *all*

¹⁶³ C.f. Dummett’s *Frege: Philosophy of Language*, Essay 13 “Can Truth be Defined”, page 454.

¹⁶⁴ The unofficial view – which I do not even pretend to defend in the main text – is that, what makes at least *some* pejorative terms *pejorative* is precisely the intention of the speaker to impart a certain use-pattern *as if* such a use-pattern can really determine an objective *meaning*. So the negative feelings are not expressed by one side of use alone, but by the self-conscious *pretention* that an incoherent set of use is in fact coherent.

To say more on this would be going too far afield. It is however fitting to point out that there is a large body of semantic literature on expressive content, content that is not properly analyzable as determinant of the truth-conditions of statements, but somehow expressed by them none-the-less. Some

Part II: Semantic Functionalism & The Problem of Laws

aspects of linguistic use are appropriately subject to the harmony requirement. Although to articulate the harmony requirement one need not be a pragmatist or verificationist about meaning, one does need to identify features of use that are relevant to meaning, rather than just *any* features of use.

Appendix to §4.1.3

“Relevant” is of course rather vague. One way the specification is satisfied is when we take the feature of use in question to partly *determine* the meaning of the expression involved. Another way is when we think that the relevant principles of use are valid *because of* the meanings of the expressions involved. More generally, for a feature of use of an expression to be relevant to meaning, it is enough that the following is true: it is not possible that the expression in question means what it means yet the feature of use does not obtain. This formulation also makes it clear that what we are after here is the *identity* of meaning, rather than the constitution of *meaningfulness*: the unabridged formulation is “relevant to the expression’s *meaning what it means*”, rather than “relevant to the ink mark’s being meaningful at all”. Despite the importance of the issue of linguistic understanding and the constitution question about meaningfulness to Dummett’s general philosophical thinking, in the context of the harmony discussion, Dummett as well as Belnap before him were concerned about the *identity* question about the meanings of logical constants.

But Dummett and Belnap also take the view that, as far as logical constants are concerned, their meanings are *determined* by the verification and pragmatic sides of their use, not the other way around. Even though this is also my view for *logical constants*, I do not think it can be generalized to a schema of identity explanation of meaning for *all* kinds of expressions.¹⁶⁵ This is reason enough for using the more general formulation of “relevance” in order to maintain the clarity that the harmony requirement itself is independent of one’s inclination or disinclination towards a particular version of use-theoretic semantics. The more directly relevant reason is that, again, I intend to apply the harmony requirement to interpret the debate about the reality of laws of nature whose participants may or may not be sympathetic to a use-theoretic semantics about law claims. The ability to formulate the

good recent examples are Potts’ “Expressive content as conventional implicature”, Kaplan’s lecture “The Meaning of ‘Ouch’ and ‘Oops’” and Kratzer’s commentary on it, or Kaufmann’s “A Modal Analysis of Expressive Meaning”. One should note, however, that these authors focus on the (compositional) modeling of expressive content and its interaction with “descriptive” content. The analysis of pejorative expression offered here is a use-theoretic analysis of *what* makes the pejorative content of such expressions expressive, not descriptive.

¹⁶⁵ For an empirical predicate, for example, the “pragmatic” side of its use does not in my view contribute to the determination of its meaning. Rather, it is determined by the predicate’s meaning.

criterion independent of use-theoretic semantic convictions will make the diagnostic application of it more convincing.

§4.1.4 Total Harmony Or Intrinsic Harmony?

The next point has been a matter of controversy in the literature. Dummett gave in his original exposition *two* precise formulations of the intuitive harmony requirement. The first formulation is what he calls “intrinsic harmony”, which is the “eliminability of local peaks”. More precisely:

...it is a requirement, for [intrinsic] harmony to obtain between introduction rules and elimination rules for [a logical constant] **c** ... that it [not be] possible, by first applying one of the introduction rules for **c** and then immediately drawing a consequence from the conclusion of that introduction rule by means of an elimination rule of which it is the major premise, to derive from the premises of the introduction rule a consequence that we could not otherwise have drawn.

– *Ibid.* pages 247/8 (underlines mine, sentences rearranged)

A “major” premise of an elimination rule for a logical constant **c** is a premise in which **c** occurs. This formulation has the double-negative form: there *cannot* be a local peak involving **c** (application of intro-rules for **c** immediately followed by application of its elim-rules) that we *cannot* level to a one-step derivation not involving **c**.

Dummett has another, *global* formulation of harmony, what he calls “total harmony”. It is basically the requirement of conservative extension: the addition of **c** to the language *minus* **c** is a conservative extension. In other words, it requires that no derivation with premises and conclusions formulated *without* **c** be made possible by the addition of **c** and its intro- and elim-rules.

The intrinsic harmony requirement for a logical constant is obviously weaker than the total harmony requirement. Total harmony implies the possibility of leveling of local peaks. Leveling of local peaks alone, however, does not guarantee that two applications, of intro- and elim-rules for a logical constant respectively, that are *not adjacent* to each other can be rearranged to form a local peak and then be leveled.

Of these two formulations, Dummett says that “the best hope for a more precise characterization of the notion of harmony lies in an adaptation of the

Part II: Semantic Functionalism & The Problem of Laws

logicians' concept of conservative extension".¹⁶⁶ Some have challenged that claim, opting for intrinsic harmony as the best formulation instead.¹⁶⁷ The issue has to do essentially with the technical fact that not all derivations of a statement with *c* as major connective will have, as its last step, an application of an introduction rule for *c*.¹⁶⁸ In other words, the option of rearranging a derivation to create a "local peak" to be leveled is not always available.

The motivation of Dummett's commitment to total harmony is complicated. It is at least connected with Dummett's wish to establish a kind of *reduction* of a logical constant's use rules to its *introduction* rules alone, or to its elimination rules alone. In other words, it is part of his attempt to establish the intelligibility of *verificationist* semantics, or *pragmatist* semantics: identifying meaning of an expression with *only* the verification side of its use, or with *only* the pragmatic side of its use. But as Dummett himself clearly recognizes, this reductionist/verificationist project is independent of the less contentious use-theoretic semantic project of formulating a harmony condition.¹⁶⁹ One *need* not assume that all elimination rules of an expression can be justified by, or are already "contained" in, its introduction rules, in order to justify the use-theoretic understanding of the meanings of logical constants, or the formulation of a condition of harmonious use.

For both the technical reasons pointed out by other philosophers as well as the fact that we do not share Dummett's *verificationist* proclivities about meaning, and so have no need to follow him in contemplating the "derivation" of the pragmatic side of use on the basis of the verification side of use, we shall stick to the *local* notion of intrinsic harmony for the discussions to follow. Moreover, since the local notion is

¹⁶⁶ C.f. Dummett's *The Logical Basis of Metaphysics*, page 217.

¹⁶⁷ A representative example is found in Stephen Read's paper *Harmony and Autonomy in Classical Logic*. One must read the paper with great caution, however, as Read does not distinguish various issues in his polemic against Dummett. He does not see, for example, that the issue of "autonomy" – complete determination of meaning by deduction rules – is completely independent of Dummett's *further*, verificationist, demand that one side of the rules (e.g. elimination rules) be "justifiable" based on the other side of rules (e.g. introduction rules). Read also confusingly formulates his critique of the so-called "Fundamental Assumption", or invertibility of rules, as directed against Dummett's notion of "intrinsic harmony". That notion is entirely unrelated to the Fundamental Assumption. In fact, Read's own proposal for a formulation of harmony is nothing but Dummett's notion of intrinsic harmony.

¹⁶⁸ This is the so-called "Fundamental Assumption" in Dummett's book.

¹⁶⁹ C.f. *The Logical Basis of Metaphysics*, page 214, "Our immediate concern is not with the question which, if either, of these aspects of our use of sentences should be taken as the central notion of the meaning-theory, but with the mere fact that linguistic practice *has* these aspects."

anyhow *weaker* than the global notion, it will be a more plausible candidate as a *necessary* condition for meaning fixation.

§4.1.5 Necessary vs. Sufficient Condition

The local harmony requirement is not, all by itself, a *sufficient* condition that a set of deduction rules in compliance with it will completely determine the meaning of a logical constant. One way the condition can fail to do that is the issue of *uniqueness*. Consider for example a close cousin of ‘tonk’, call it ‘tink’. ‘Tink’ has the introduction rule of ‘and’ and the elimination rule of ‘or’. Formally:

$$\begin{array}{c}
 \begin{array}{ccc}
 & [A]_1 & [B]_1 \\
 & \ddot{C} & \ddot{C} \\
 A \text{ tink } B & & \\
 \hline
 & C & \\
 & & 1 \text{ (Tink-} \mathbf{Elimination} \mathbf{)} \\
 & C &
 \end{array} \\
 \\
 \begin{array}{ccc}
 A & & B \\
 \hline
 & A \text{ tink } B & \\
 & & \text{(Tink-} \mathbf{Introduction} \mathbf{)}
 \end{array}
 \end{array}$$

Obviously, “tink” has an introduction rule that is ‘too strong’ for its elimination rule. Consider an object language obtained by adding ‘tink’ to the ordinary English language. We can interpret ‘tink’ either of *two* ways, either along with its strong elimination rule, or along with its weak introduction rule. The first says that ‘tink’ means *and*, the second says that ‘tink’ means *or*. Moreover, both interpretations interpret the rest of the ‘tink’-enriched language in the ordinary, homophonic way. Clearly these are distinct interpretations, since ‘or’ and ‘and’ are not synonymous in our metalanguage. Moreover, both interpretations are compatible with the use-rules of ‘tink’ in the sense that both the elim-rule and the intro-rule specified above will, under either of these interpretations, remain good inference rules.

There are philosophers who use “harmony” to mean a more stringent condition that is also sufficient for *fixing* meaning. They do that, for example, by including a version of Belnap’s uniqueness condition I mentioned earlier. Apart from a very limited case that I shall turn to presently, this is not what I shall do in this essay. The reason, again, has mainly to do with the *diagnostic* use to which I shall put the harmony requirement. On the one hand, since all that is needed is an explanation why the concept of nomological connection has been felt to be *in-coherent*, the

Part II: Semantic Functionalism & The Problem of Laws

harmony requirement does not *need* to be formulated so strongly to do the diagnostic work. On the other hand, the diagnosis will be more persuasive if the harmony requirement is not formulated in such a way as is acceptable only to use-theoretic semanticist. What distinguishes a use-theoretic semanticist from her colleagues, however, is precisely that she insists that the use-rules governing, for example, a logical constant, are *sufficient* to determine its meaning, as long as they conform to some strong version of the harmony condition.

What we have now arrived at, then, is a *necessary* condition for an expression's expressing a coherent concept. The condition is a *local* one, and demands that the *verification side* of the expression's use be in harmony with the *pragmatic side* of its use, where being in harmony means that the pragmatic principles do not allow one to draw consequences from a statement that are unwarranted by the grounds deemed adequate for that statement by the verification principles.

§4.1.6 Presumption of Unique Meaning

I said that there is a small exception to my *general* refusal to formulate a harmony condition strong enough to guarantee the *uniqueness* of meaning. We also saw above – with the example of 'tink' – that the harmony condition I have adopted is generally not sufficient for fixing a unique meaning. The exception that I have in mind is a very special sort of situation in which satisfaction of the weak harmony condition I have spelled out suffices to secure a *presumption of unique meaning*. This is the case where the verification side and pragmatic side of use are *practically identical*. Let me illustrate with a simple example.

Brandom claims that the assertion of a conditional is warranted when a certain corresponding inference is good, and accepting the conditional has the consequence that one must also accept this corresponding inference. In Brandom's account, therefore, a conditional of a certain kind has, in effect, a condition of warrant that is *identical* with its consequence: both are (the goodness/acceptance of) a certain inference. Assuming that there *is* a kind of conditional that is used the way that Brandom describes, the near identity of consequences and conditions of warrant means that we must *presume that the conditional has a determinate meaning*.

I will show that, in the case of (certain primitive variants of) law statements, we have an analogous situation: the condition of warrant is of the same strength as the

consequence. This compliments the harmony condition because it also guarantees, beyond harmony, a presumption of *uniqueness of meaning*.

§4.2 The Problem of Laws of Nature

§4.2.1 Overview

While Hume talked about *evidence* for necessary connections between matters of fact, more recent, post-linguistic-turn literature tends to talk about features of *statements* of laws of nature. The debates about whether there is a distinctive aspect of *reality* – namely law-like connections – now take the form of the question: do law claims express a distinctive kind of *content*? The point of this section is to show that the fact that there is space for debate at all – a space for developing skeptical dialectics and refutations and so on – is largely due to a certain feature of the constellation of properties that *both* the realist and the antirealist attribute to statements of laws of nature. The feature in question, as the reader has already been primed to anticipate, is of course the *disharmony of use*. Given this goal, our discussion must begin with a review of this constellation of properties of law-statements that enjoy almost universal, unquestioned approval.

The following list should include all the major claims or assumptions made in recent debates about the reality of laws. The items on the list are such that, as van Fraassen put it nicely, none is entirely undisputed, but all are generally respected.

- (a) A statement of law purports to express a *universal* truth;
- (b) The descriptive general terms appearing in a law statement occur in *opaque* positions; in other words, the law-operator is an *intensional* operator;
- (c) A statement of law can be confirmed by another statement of law from which it *follows deductively*;
- (d) A statement of law can be confirmed by its *observed instances*;
- (e) A statement of law supports a certain kind of counterfactuals;
- (f) Statements of law tell us what *must* happen, as opposed to merely what has happened or will happen;
- (g) A statement of law enable us to *predict* the future;

Part II: Semantic Functionalism & The Problem of Laws

- (h) A statement of law can be used to *explain* phenomena that fall under it as instances;

Before going through these in details, some general remarks are in order. These eight features are not of equal significance. The first one, for example, is not *distinctive* of law statements. All universally quantified claims purport to express universal truths, as well. The rest seven features are usually mentioned to distinguish a law statement from a merely universally quantified one. The second feature (b), like the first, is a *strictly semantic* observation, as opposed to an observation about how law claims are *used* by us, for example.

The remaining six features of law statements are about the *use* we make of law statements. Of these, (c) and (d) are plausibly classified as *verification principles* in the sense of Dummett discussed earlier: they describe what can count as giving warrant to a law claim or conclusively establishing what it expresses. The next three, (e), (f), and (g), belong to the *pragmatic* side of use likewise discussed above. According to (e), accepting a law claim has the consequence that one also has to accept a certain type of counterfactual conditionals. (f) and (g) describe two other sorts of consequences of statements of law: modal statements and statements about the future. The last one, (h), is indeed also a feature of *use*, but does not seem to be either on the verification side, or on the pragmatic side. In fact, it might even seem that the explanatory use we make of laws is not *essential to the meaning* of law statements. For, even if we *did not* use law claims for explanations – but retained the other uses – law statements would not have a different meaning. While it is important for the discussion below to note that (c) and (d) fall on the verification side of use, while (e)-(g) on the pragmatic side of use, I shall not insist on the status of (h).

What makes (h), or the explanatory role of law statements, especially significant is that it has been seen as probably *the* central argument for law claims' having a content going beyond universally quantified claims. I shall not get into the debate about this argument and various fancy gymnastics performed on the notion of explanation to avoid its conclusions. My concern is to show how a structural feature of the verification side and pragmatic side of use for law statements, namely, a structural feature characterizing the set of five items (c)-(g), renders a skeptical dialectic both inevitable as well as a resolution unlikely. In the rest of the section I shall go through these supposed features of law claims, while tracing a picture of

inexorable tension between the realists and the anti-realists arising from items in that list.

§4.2.2 Intensionality of the Law Operator and Two Types of Accounts

There are two sorts of accounts of law statements. To describe them, let us agree on some terminology. In the statement of law “It is a law of nature that copper conducts electricity”, let us call “It is a law of nature that” the *law-operator*, and the rest of the statement “copper conducts electricity” the (*substantive*) *core* of the statement. The substantive core of a law statement is always some general, perhaps generic, statement. According to what I shall call a (*strictly*) *semantic* account, the application of the law-operator to a general statement *transforms the content* of that statement into a content of a special type. This contrasts with what I shall call a *pragmatic* account, which says that adjoining the law-operator to a general statement does not change its content, but adds to it various claims about how the content of the original general statement is *used* in special ways.

A classical example of the pragmatic sort is what some have called the Mill-Ramsey-Lewis account. Ramsey takes laws to be “consequences of those propositions which we should take as axioms if we knew everything and organized it as simply as possible in a deductive system”.¹⁷⁰ Lewis refines that to read “...contingent generalisation is a law of nature if and only if it appears as a theorem (or axiom) in each of the true deductive systems that achieves a best combination of simplicity and strength”.¹⁷¹ To say that *it is a law of nature* that copper conducts electricity, according to Lewis, is to say that (i) copper conducts electricity, and (ii) this statement appears as a theorem in every deductive system that has “achieves a best combination of simplicity and strength”. The additional information we convey by calling a universal truth a law, namely (ii), is its position in a part of our linguistic practice, the part where we try to formulate true deductive systems to capture our observations about the world.

One kind of semantic account has it that adding “It is a law of nature that” to “copper conducts electricity” transforms a universally quantified claim into a claim of strict implication. Another type of semantic account that has been popular since the

¹⁷⁰ Ramsey 1929, page 150.

¹⁷¹ Lewis 1973, page 73.

Part II: Semantic Functionalism & The Problem of Laws

late seventies says that the effect of the law-operator is to turn a universal quantification, a claim about concrete *objects*, into a claim about a relation between *universals* these objects instantiate. In Dretske's words, the law-operator effects an "ontological ascent" in the content of a general claim.

One nice way to think about these two types of accounts of laws, as Dretske's similar discussion in his seminal paper on laws of nature suggests,¹⁷² is in terms of the feature (b) listed above: namely the opacity of the law-operator. What (b) says is that, amongst the following three statements:

- (A) Copper conducts electricity.
- (B) It is universally true that copper conducts electricity.
- (C) It is a law of nature that copper conducts electricity.

only in the last statement the general terms "copper" and "conducts electricity" occupy opaque positions, that is, positions that do not allow intersubstitution of *co-extensional general terms* without altering the truth-value of the whole statement. Suppose, for example, that copper is the material for tools most widely used by humans before 3000BC. So "copper" will be co-extensional with "the material for tools most widely used by humans before 3000BC". While it is still a universal truth that the material for tools most widely used by humans before 3000BC conducts electricity, it is highly unlikely that it is a *law of nature* that the material for tools most widely used by humans before 3000BC conducts electricity.

The semantic account of law statements accounts for the opacity (C) by a supposed transformation of content effected by the law-operator. The pragmatic account explains the opacity through the opacity of the *additional* content added by the law-operator: the additional content states how the statement "copper conducts electricity" is used and so *mentions* but does not use it; but mentioning a statement creates an opaque context for the general terms occurring in it.

¹⁷² C.f. Dretske's *Laws of Nature*. Dretske contrasts anti-realist views with his as two possible ways of meeting the opacity requirement, without drawing a broad distinction based on that requirement. Moreover, Dretske's contrast is formulated in terms of treating the difference between laws and universally quantified claims as an *intrinsic* difference vs. as a *functional* difference. We cannot quite adopt that formulation given that differences in *content* might *also* be a matter of use-functional difference. A way to preserve the spirit of Dretske's contrast is to distinguish functions relevant to *content* from functions relevant to the *pragmatics* of a statement. But if we have to do that, it is more straightforward to use the content vs. pragmatics distinction directly, without the detour through functions.

While pragmatic accounts of law statements do not need to entail that there is no *fact* of the matter which statements are laws, they do deny there are *nomological* facts in the *world*, whose obtaining is independent of *our* constitutions and parochial practices. What motivates this anti-realist attitude?

§4.2.3 Apparent Disharmony I – The Problem of Induction

To simplify discussion, let me give names to the verification and pragmatic principles of use identified above for the law-operator:

Verification Principles:

- (c) justification by deduction rule;
- (d) justification by observation rule;

Pragmatic Principles:

- (e) counterfactual consequence rule;
- (f) modal consequence rule;
- (g) prediction rule;

Let us suppose for the following few paragraphs that the properties we have identified as the verification and pragmatic sides of use for law statements are *relevant* to the meaning of law statements and the law constant. I call the disharmonies I am going to identify *apparent* because they obtain only if we make this meaning-relevance assumption. As we shall see, philosophers *respond* to apparent disharmonies partly by denying that some of these principles (c)-(g) are relevant to meaning. But the pre-theoretic, dialectically initial position is that they are.

Under this assumption then, there is an outright disharmony between the *observation rule*, and the *prediction rule* for the law-operator. An observed instance for the law statement that copper conducts electricity is presumably a piece of copper observed to conduct electricity. Since anything observed has already happened, the copper-conducting-electricity instances that have been observed are all events in the *past*. Predictions licensed by the law claim, however, concern the *future* behavior of pieces of copper some of which no one has ever yet seen. So, the observation rule (d) recognizes observation reports of events in the past as sufficient ground for something that, according to the predication rule (g), licenses conclusions about events in the future. The consequence allowed seems obviously to go beyond what is recognized as adequate grounds.

Part II: Semantic Functionalism & The Problem of Laws

To put the matter another way, if we apply the prediction rule (g) – which eliminates the law-operator – immediately after the observation rule (d) – which introduces the law-operator – we shall not be able to “level” the local peak to obtain a *direct* derivation of a future prediction from observation sentences about past events. From claims like:

X1 is a piece of copper observed to be subject to electric potential difference at T1, and to carry electric current at T1;

X2 is a piece of copper observed to be subject to electric potential difference at T2, and to carry electric current at T2;

...

Xn is a piece of copper observed to be subject to electric potential difference at Tn, and to carry electric current at Tn;

X1, ... Xn are *all* the pieces of copper observed to be subject to electric potential difference;

we cannot derive:

If Y is a piece of copper subject to electric potential difference at T, then it *will* carry electric current at T.

(In these formulations I have elaborated the dispositional predicate “conducts electricity” into “carries electric current when subject to electric potential difference”. An instance of the law is a piece of copper satisfying the *conjunction* of two properties: being subject to electric potential difference *and* simultaneously carrying electric current.) Note that the *pattern* of this inference is obviously not valid, which we can see by substituting “is/was in my pocket” for “conducts electricity” in the law claim.¹⁷³ But if we are allowed to use the prediction and observation rules as well as

¹⁷³ Doing that requires that, correspondingly, the complicated predicate in the description of instances “was subject to ... and carried ...” be replaced with “is/was in my pocket” also. The difference between instance-description predicate and predicate in the law-claim results from the fact that the example I am using contains a *dispositional* predicate: conducts electricity. Whenever this is the case, the universality of the law is *partly hidden in the dispositional predicate*, rather than merely consisting in the fact that the law claim covers all individuals falling under these predicates. In the present case, the instances are not just pieces of coppers – which are ostensibly the individuals falling under the predicates “copper” and “conducts electricity” – but instances of a piece of copper *being subject to electric potential*.

More formally, if the law claim “L(G(A, B))” contains a dispositional predicate “B”, which has the meaning “when ϕ s then ψ s”, then the law claim can be unpacked as: L(G(A & ϕ , ψ)). I could

the law-operator, it appears that we *can* make the inference, mediated by the law claim:

It is a law of nature that copper conducts electricity.

But might it not be that the use of the law-operator *renders explicit* what is hidden in the observational sentences? Now *what* could be hidden in the observation sentences, which are singular claims of the form “X ϕ -ed”? The prediction sentence and observation premises – excluding the exhaustive clause – have the same structure, use the same general terms. They differ only in that they talk about different *particulars* and different *times*. So it is not like there could be hidden conceptual connections between the general terms employed in the premises and the conclusion, the way there is a connection between “X is a *bachelor*” and “X is *male*”. Is the hidden connection between the *particulars* talked about then? But what kind of connection is there between pieces of copper at times in the past and a piece of copper at a future time? In any case not enough a connection to warrant a similar inference with “is/was in my pocket” in place of “conduct(ed) electricity”. In short, there is little case to be made that the use of the law-operator is nonintrusive and merely makes explicit hidden connections – either of a conceptual sort or between particulars – between the observation sentences and the prediction sentence.

What we have then, is a situation that is structurally *exactly* like the problem with “tonk” or “Boche” which we have seen before. The mediation of “tonk” and its use-rules allows us to derive from “p” a claim “q”, for arbitrary “p” and “q”, which is blatantly bad. The use of “Boche” allows us to infer, equally without validity, from being German in nationality to being prone to cruelty. We saw what was plausible to say about “tonk”, given its disharmony of use: its use-rules do not determine a coherent concept. In the same way, if the observation rule (d) and the predication rule (g) are part of the use of the law-operator that is relevant to its meaning, then we must say, it appears, that the law-operator expresses no coherent concept. There is, in other words, no *meaningful* expression whose meaning allows for use-rules (d) and (g).¹⁷⁴

have used a different example in the main text, say, the law of the ideal gas, which involves only measurement-predicates that are not obviously dispositional. In any case, this complication does not affect the point being made in the main text.

¹⁷⁴ There is another response, that of Hume’s, which corresponds to what I have said about “Boche”. See more on this in §4.3.1.1.

Part II: Semantic Functionalism & The Problem of Laws

The tension between (d) and (g) is only *one* of several apparent disharmonies between the verification side of use and the pragmatic side of use for the law-operator. It is the apparent disharmony that underlies the so-called *Problem of Induction*. To *solve* the Problem of Induction is to resolve the tension between the prediction rule and the observation rule just sketched above. I shall not delve into the complicated and sometimes ingenious dialectics developed to do that. Instead, let me turn to *another* apparent disharmony, which, for lack of a better name, I shall call the *Problem of Counterfactuals*.

§4.2.4 Apparent Disharmony II – The Problem of Counterfactuals

The Problem of Counterfactuals is the apparent disharmony between (d) on the one hand, and the modal and counterfactual rules ((e) and (f)) on the other hand. The other rule on the verification side – deduction rule (c) – is not important in the realist-vs-anti-realist debate.

The reason that I use the phrase “Problem of Counterfactuals” to refer to the tensions between *both* the counterfactual *as well as* the modal rule on the one hand and the observation rule on the other, is that the cash value of modal claims in our current context more or less are the counterfactuals that follow from them. To explain, we need to spell out the modal and counterfactual rules first. In one version of the modal rule, we can make the following derivation:

$$\frac{\begin{array}{l} L(G(A, B)) \\ A(N) \end{array}}{\Box B(N)} \text{ (L-Modal Elimination-}\alpha\text{)}$$

Here ‘L’ stands for the law-operator. ‘A’ and ‘B’ are general terms, and ‘G(A, B)’ is the general/generic statement that forms the substantive core of the law statement. ‘N’ is some singular term. Another version of the rule, which is at least as strong, is:

$$\frac{L(G(A, B))}{\Box (A(N) \rightarrow B(N))} \text{ (L-Modal Elimination-}\beta\text{)}$$

Finally, a yet stronger version of the modal rule is:

$$\frac{L(G(A, B))}{\Box (A(N) \rightarrow B(N))} \text{ (L-Modal Elimination-}\gamma\text{)}$$

$$\Box \forall x(A(x) \rightarrow B(x))$$

The counterfactual rule can be similarly represented formally. If we use “ $\Box_c \rightarrow$ ” for the counterfactual conditional forming device, we can represent the rule this way:

$$\frac{L(G(A, B))}{A(N) \Box_c \rightarrow B(N)} \text{ (L-Counterfactual Elimination)}$$

The sense in which the modal rule’s cash value lies in the counterfactual rule is that, first, the modal conclusions of the β - and γ -versions of the modal rule imply the conclusion of the counterfactual rule, in cases where the presupposition of the counterfactual is satisfied: i.e. where $A(N)$ is false. Second, and more importantly, counterfactuals form our epistemic access to the modal claims. So, to illustrate the falsity of the alethic modal claim, to use Dretske’s example, that, necessarily, dogs born at sea are cocker spaniel cubs, we point to the falsity of the counterfactual: if my dachshund had been born at sea, she would have been a cocker spaniel. This is how we *know* and *demonstrate* that the modal claim is true or false. Our intuitions are stronger, or appear to work more directly, in the case of counterfactuals than for modal claims. For these reasons, I shall limit myself in the following exposition on the tension between the counterfactual rule and the observation rule.

The Problem of Counterfactuals is simply the fact that the ground recognized by the observation rule as adequate for a law claim – which are reports about instances falling under the law claim – does not warrant the conclusion that the counterfactual rule allows us to draw from the law claim. Again, a classic case of disharmony. Even if we have reports confirming *every* actual dog born at sea to be a cocker spaniel – ignoring for the sake of the argument that we cannot in fact observe all such instances, certainly not those that occur in the future – we are still not warranted to draw the conclusion that, *had* my dachshund been born at sea, it *would have been* a cocker spaniel. The premises are conceivably all true – it is conceivable that *in fact* all dogs born at sea are cocker spaniels – but the counterfactual conclusion is plainly false: my dachshund would have been a dachshund no matter where it had been born.

§4.3 Realism vs. Anti-Realism About Laws of Nature

Armed with the two patterns of apparent disharmony of use, we are now ready to give a diagnosis of some of the main positions and arguments in the debates about the reality of laws of nature.

§4.3.1 The Anti-Realists

§4.3.1.1 Hume

The first of the four types of positions is that of Hume, who was also the first philosopher to formulate one of the disharmonies into a skeptic argument against “necessary connections”: namely, he raised to everlasting prominence the Problem of Induction.

In what I shall call *Hume’s signature argument*, he concentrates on the transition from *past* observed “constant conjunction” of two or more types of events to *future instances* of these events, the transition known as *induction*. The argument is that induction cannot be justified except with the help of a *principle of the uniformity of nature*.¹⁷⁵ The latter, however, cannot be established deductively¹⁷⁶; but nor can it be justified “by probability” on pain of circularity. This is because “probable reasoning” requires a form of justification by induction,¹⁷⁷ which itself is supported by the very principle of uniformity to be established. Hume’s signature argument has the form of an *epistemological argument*, but it is clearly just another formulation of the disharmony between the observation rule and prediction rule we have just identified. The fact that we described these use-rules for *law* statements and that Hume talks about “constant conjunctions” instead, does not vitiate the point. For the disharmony persists if there is *any* universal statement – whether it be about laws or

¹⁷⁵ Hume’s formulation of this principle is “that instances, of which we have had no experience, must resemble those, of which we have had experience, and that the course of nature continues always uniformly the same”. (Page 89, *A Treatise Of Human Nature*, Part III, Section VI “Of the inference from the impression to the idea”).

¹⁷⁶ Or in Hume’s words, “by demonstration”.

¹⁷⁷ Hume argues for this claim beginning with the last full paragraph on page 89, *ibid*, until the next page. The argument has two steps: (i) “’Tis ... necessary, that in all probable reasonings there be something present to the mind, either seen or remember’d; and that from this we infer something connected with it, which is not seen or remember’d”; and (ii) “The only connexion or relation of objects, which can lead us beyond the immediate impressions of our memory and senses, is that of causes and effect”. (Underlines added) The relation of cause of effect is established through induction. Therefore, “probable reasoning” requires a form of induction.

“constant conjunctions” – that can be established by past observations alone and yet can license predictions about the future. The disharmony arises from the fact that both observation and prediction rules are rules for the *same* expression, and is independent of what that expression is purportedly about. The epistemological formulation of Hume’s signature argument gives a special emphasis on the *inadequacy* of the verification side of use (rather than, say, the largesse of the pragmatic side of use), whereas our identification of disharmony is neutral.¹⁷⁸

Hume had other arguments, which are again couched in epistemological terms, that is, they emphasize the inadequacy of the verification side of use. Hume argued, for example, that the idea of “necessary connexions”, “power”, or “causal efficacy” are not adequately grounded in experience: we do not perceive them directly, nor does the observation of “constant conjunction” give rise to such ideas that go beyond mere conjunction of events. In the way we have set up things, some of these arguments pertain directly not to general law claims, but to *singular* claims. The right way to deal with them is as Anscombe did,¹⁷⁹ by pointing out that Hume’s claim that we can only perceive so much is wrong intuitively, and understood as a technical claim simply begs the question.

Returning to Hume’s signature argument, the *lesson* that Hume drew from it is that the pragmatic side of use, namely, the prediction rule, is *not essential to meaning*. He explains the inference to future instances in terms of the *habits of the mind*. This is like resolving the disharmony involved in the use of “Boche” by declaring that the pragmatic side of use – the inference from being a “Boche” to being prone to cruelty – is merely *expressive* of the speaker’s *state of mind*, rather than warranted by the meaning of “Boche”, which is the same as the meaning of “German”.

§4.3.1.2 Contemporary Anti-Realists

Contemporary anti-realists about laws of nature are not straightforward Humeans. More than anything else, what distinguishes most of them – as well as the realists – from Hume is their focus on the *Problem of Counterfactuals* as opposed to the Problem of Induction. This formulation is paradoxical. For a vast amount of recent literature on laws of nature is on how to “solve” the Problem of Induction. What I

¹⁷⁸ For the possible objection that the observation rule is not relevant to meaning, see the end of §4.3.2.1 below.

¹⁷⁹ In her Inaugural Lecture at Cambridge University, in 1979.

Part II: Semantic Functionalism & The Problem of Laws

mean is this: in contemporary debate, the Problem of Induction has been *accepted*, by all sides, as a problem *to be solved*. That does not mean that anyone would argue that the inference from past instances to future instances is philosophically unproblematic. Rather, the Problem of Induction is regarded as a *technical problem*, rather than a problem of principles. So, the anti-realists do not work with universal claims about *so far observed* instances. They work instead with universal claims about *all* instances, including those in the past, present, and future. On that basis, the anti-realists argue that above and beyond such regularities of *unlimited* scope, there are no law-like connections.

Consider, for example, the following passage from D.H. Mellor, which is typical in its acknowledgement of the Problem of Induction as a bipartisan problem:

Now if giving laws one content rather than another made the problem of induction soluble for them, this would be a strong argument for giving them that content. But since I believe no such solution is presently available for any credible content, I must look to other arguments. Hume's problem does, however, provide a reason for preferring weak readings of natural laws. The less a law says, the less there is to be certified in claiming it to be true.

– *Necessities and Universals in Natural Laws*, pages 848-9

Relegating the Problem of Induction to the “technicalities” has the misleading effect that what is at issue appears to be “merely” an epistemological problem, not a conceptual one. That is, the inadequacy of observed instances is perceived *not* as undermining the coherence of the very concept of a universal claim with unlimited scope establishable by these observed instances, but as a problem whose articulation *presupposes* that there *is* that sort of universal content: given that law claims cover instances both in the past and the future, how are we ever *warranted* to make such claims? As we shall see later, this merely disguises the conceptual problem.

In any case, the epistemization of the Problem of Induction still leaves the other disharmony, the Problem of Counterfactuals, untouched. We have seen that Hume's way to dissolve the first kind of disharmony is by declaring that the prediction rule is not relevant to meaning. The modern anti-realists have a similar strategy: they *discredit the pragmatic side of use*. Their favorite way of dissolving the second kind of disharmony is to say that the modal and counterfactual rules, which are the pragmatic use rules that stand in tension with the observation rule, are *not*

what they seem. For, according to the anti-realists, modal and counterfactual statements do not really have objective truth-values. One of the clearest examples of this strategy can be found in van Fraassen's writings, where he spends much energy to give an account of counterfactuals according to which they are context-sensitive in such a way that they are not true or false objectively, but depending on "which regularities we decide to raise to the status of a law" and therefore to keep invariant under the counterfactual condition. Van Fraassen's predilections for context-sensitivity as an argumentative strategy is partly a product of the time: there was an exploding amount of research in semantics beginning in the seventies on indexicality, for example. But van Fraassen's application of the context-sensitivity strategy is less successful in the case of counterfactuals as in the case of scientific explanation. I shall not get into a detailed appraisal of it here, except to point out two things. First, the untenability of non-factualism about counterfactuals will become clear later when I introduce the notion of *desiderative inferences*; for, although the intended meaning of counterfactuals can be hard to get one's finger on due to the richness and varieties of natural language, whether an inference pattern in a *practical deliberation* is good or not is clearly an objective matter. Second, van Fraassen's non-factualism about counterfactuals independently gets him into a problem due to his claim that *observability*, which is closely related to counterfactuals, is an objective matter.¹⁸⁰

The take-away now though is that the anti-realists resort to the general strategy of discrediting the pragmatic side of use to disarm a disharmony. We should also note that there may be inconsistency in the intermediate position occupied by most contemporary anti-realists about laws, namely the position between a strict Humean denial of universal content with *unlimited* scope extending into the future on the one hand, and a fully realist acceptance of a law-like content on the other. The reason is that claims about the future may be *modal* in nature. To say that it will rain tomorrow, intuitively, is to say that *necessarily*, it will rain tomorrow.¹⁸¹ If this is so, by virtue of admitting universal content with unlimited scope, the anti-realist already admits modal content of *some* sort. That makes it puzzling what would then motivate

¹⁸⁰ For this aspect of van Fraassen's view, c.f. the exchange between van Fraassen and James Ladyman, especially Ladyman's *Constructive Empiricism and Modal Metaphysics: A Reply to Monton and Van Fraassen*.

¹⁸¹ Prominent supporters of this modal view about simple future claims include C.S. Peirce and A. N. Prior. For an exposition, see Prior's *The Formalities of Omniscience*, especially the last third of the paper, beginning with page 124.

the denial that counterfactual consequences of law statements are not objectively true or false. I suspect the reason is that these anti-realists assume an *ahistorical* view about the content of (at least scientific) knowledge. But this, like countless many other issues in the debate, will be left aside for the rest of the essay.

§4.3.2 The Realists

§4.3.2.1 *The General Position*

Whereas the anti-realists discredit the pragmatic side of use, the realists *harp* on it. So they argue that, if law statements really just express universally quantified claims with unlimited scope plus some semantically irrelevant features of use, how do they justify modal and counterfactual conclusions we draw from them but cannot draw from universally quantified claims? Another favorite argument by the realists is, as already mentioned, the argument from the explanatory power of law statements. Sometimes this is thought as related to the arguments from modal and counterfactual consequence: if a law statement implies the *necessary* co-occurrence of two events, it implies something that is not the same as *mere* co-occurrence, so can genuinely explain the latter. Respecting the genuine modal character of law statements allows one to explain their explanatory power. But if laws were really essentially universally quantified truths, that power remains a mystery.¹⁸²

Besides harping on the counterfactual rule and explanation, there is another argument frequently advanced by the realist that is, curiously, formulated as an *epistemological* one. It is curious because, as we have seen, it is precisely the *opponents* of the realists – namely the (strict) Humeans and anti-realists – who tend to harp on the verification side of use while simultaneously discrediting the pragmatic side of use. This argument from the realists is that, one cannot confirm a universally quantified claim by checking only *some* of its instances: the probability that the *unchecked* instances fall under the claim is not increased by the fact that the ones *so far checked* do fall under it. In contrast, for law statements:

Our confidence in them increases at a much more rapid rate than does the ratio of favorable examined cases to total number of cases. Hence, we reach the point

¹⁸² For a good exposition of this realist's argument from explanation, see Alexander Bird's *Laws and Criteria*, section III.2, pages 525-529.

of confidently using them to project the outcome of unexamined situations while there is still a substantial number of unexamined situations to project.

– Dretske, *Laws of Nature*, page 256.

The idea is that, we cannot solve the problem of *projecting* from observed instances to unobserved instances – essentially the problem of induction – unless we assume that the projection is *mediated by a strictly stronger content* than universal quantification. That is, unless we assume that we *first* go from observed instances to *laws* – something genuinely stronger than universally quantified claims – and *then* from laws to unobserved instances, the problem of projection or induction will remain unsolvable.¹⁸³ Of course, the realists do not themselves offer an account of *how* the confirmation of a law claim from observed instances goes. They just think admitting stronger contents gives them a better chance to solve the induction problem. The formulation of this argument in terms of the induction problem is explicit in a recent exposition by Alexander Bird:

If there is no connection between instances of a regularity, then there can be no reason to infer from certain of its instances to other instances. ... On the other hand, if there is a connection between the instances of a regularity, i.e. they all instantiate the universal or property F which itself has the property of bringing about the co-presence of the property G, then we do have a justification for inferring from the observation of Fs being G that all Fs are G. The point of these remarks is not to provide a solution to either of the two problems of induction [i.e. one to universally quantified contents, the other to law contents]. However, they provide some insight into the nature of induction, that the inductive schemas given above [i.e. “All observed Fs are Gs therefore All Fs are Gs”, and “All observed Fs are Gs therefore the next F is G”] are justified only if something like the following is justified too:

¹⁸³ This is the gist of the well-cited feature that “laws can be confirmed by its instances”. The paradoxical air of it as an argument against the anti-realists results from the fact that it can be misleadingly read as argument on the level of the *object practice*, when it really is an argument on the *meta-level*, about how best to *analyze* the object practice. So the argument can appear to issue in the recommendation that, if we have difficulty confirming a universal claim from its observed instances, we should first formulate a much stronger corresponding law claim, which *will* be confirmable by the observed instances. But that is not how the argument is to be understood. A nice formulation of the argument that is not misleading in this way can be found in Alexander Bird’s *Laws and Criteria*. See the main text below for details.

Part II: Semantic Functionalism & The Problem of Laws

All observed Fs are Gs therefore it is a law that Fs are Gs
where, furthermore, the law that Fs are Gs is more than just a regularity.

- Bird, *Laws and Criteria*, pages 530-531. (Underlines mine)

In other words, induction to universally quantified content cannot be justified unless “induction” to a stronger type of law contents can be justified.

Apart from this indirect argument just described, the realists appear to simply *ignore* the verification side of use for the law-operator. So, they do not attempt as hard to tell us how we become *warranted* to assert modal contents by observing instances, as they insist on the indispensability of modal content if law statements are to have counterfactual consequences and to explain. This is of course what one would expect given that the realists want to endow law statements with a type of content closer to their modal/counterfactual consequences than to their observation warrants. But at least part of the reason is that their opponents, the more recent anti-realists, admit they have a problem with the verification side as well. Yet there is an asymmetry between them. Whereas Hume and the anti-realists can dissolve the disharmonies by discrediting the pragmatic side of use – either by saying that particular pragmatic rules are not relevant to meaning or by saying that they are not what they seem – the realists cannot adopt the corresponding strategy by discrediting the verification side of use. The reason is simple: if a realist is to claim that the observation rule is not relevant to the meaning law statements, then he must propose *some* verification principle that *is* so relevant. If he does not, the realist must countenance statements for which there is nothing that would even *count*, in view of the meanings of these statements, as providing warrant for them. Alternatively, if he claims that the observation rule is not what it seems, he will have to tell us what it *really* is about. But what could that be, if it is not to be just that, *observation of particular instances?* .

What all of this amounts to is that the realists have no strategy available to them to *discredit* the verification side, and so they simply *ignore* it. But that means they have an epistemological problem, which, though they partly share with their anti-realist opponents, is somewhat more severe for them because of their belief in a stronger content than universally quantified ones.

§4.3.2.2 *The Criterial View – Transition To A Use-Theoretic Account*

There is one strand in the realists' thinking that is worth a separate mention. The (strict) Humeans, the anti-realists, as well as the realists discussed so far, when faced with disharmony, all cling onto one or the other of the two horns and either discredit or ignore the other horn. The criterial view, however, attempts, quite courageously, to keep both discordant horns. They are *realists* in that they treat the modal and counterfactual rules seriously. And, rather than ignoring the verification side of use, the criterialists *insist* on the relevance of the observation rule to the meaning of law statements.

The other feature of the criterial view that distinguishes it from the rest of the pack is that, unlike the realists and anti-realists, it explicitly avoids giving a *truth-conditional analysis* of law statements, opting for some variant of functionalist semantics instead. We shall see why this is the right approach presently, when we consider below the various difficulties truth-conditional approach has. To see how a criterial view attempts to finesse the problems for law statements, let me turn now to the most detailed exposition of it, in Alexander Bird's *Laws and Criteria*.

The idea of a criterion for employing a certain expression is, according to Bird, the idea of *a priori* evidence. An often cited wittgensteinian example is that someone's eating with gusto is a criterion for saying that he is hungry. The relation between being hungry and eating with gusto is not like the relation between a footprint of a certain shape and the presence of a tiger nearby. In order to know that a certain kind of footprint is evidence that a tiger is roaming nearby, one has to have *investigated* or otherwise *learned* something empirical about tigers. Not knowing that does not discount me from grasping the *concept* of a tiger: I just do not know the distinct shape of a tiger's footprint. But not knowing that someone's eating with gusto is evidence for his being hungry *would* cast doubt on whether I know what it *means* to say of someone that he is "hungry". Learning the former, criterial fact is part of learning the meaning of "hungry".¹⁸⁴

The "meaning constituted by use" aspect is just one feature of the criterial view, a feature that it shares with *any* attempt to construct an identity-explanation of meaning on the basis of linguistic use. The criterial view is furthermore an element of

¹⁸⁴ Bird adopts a formulation from Christopher Peakcocke that is even stronger. He characterizes a "criterial concept" to be that concept to possess which it is *necessary and sufficient* that the speaker knows the criterial facts about it. See *Laws and Criteria*, page 515.

Part II: Semantic Functionalism & The Problem of Laws

a subspecies of use-functionalism, the so-called “anti-realist” semantic view. That view has the *additional* commitment that the kind of use that constitutes meaning is *exclusively on the verification side*. What makes a semantic view an anti-realist semantic one is moreover the fact that the “verification side” is interpreted strictly internalistically, in that what gives warrant for making an assertion employing the “criterial concept” is the *recognition* of a state of affairs – that a criterion is fulfilled – rather than a state of affairs that the speaker himself need not be aware of.

§4.3.2.2.1 Defeasibility and Ontological Gap

Yet *the* distinguishing feature of the criterial view, at least of Bird’s variety, is its attempt to bridge what is to function as criterion on the one hand and the claim it is supposed to be a criterion for on the other hand, by the idea that a criterion is *defeasible*. Though the criterial view originates from interpretations of later Wittgenstein – it is however disputed by some that Wittgenstein himself held such a view – it is in its substance an adaptation of the anti-realist semantics for logical vocabulary to non-logical vocabulary. For logical vocabulary, what plays the role of “criteria” are just proofs that *conclusively* establish a claim. For some non-logical vocabulary, there might be similar conclusive conceptual links between warrant and claim, for example between being male and unmarried on the one hand and being a bachelor on the other. But the application of the anti-realist semantic strategy is thought to be fruitful precisely in those cases where a serious gap between available *evidence* and the content of a claim is perceived. This is why its first “application” was to the problem of other minds: it was thought that a devastating gap exists between the content of the claim that there are other minds and the behavioral evidence for that claim. This gap is usually conceived, not in terms of *disharmony* of two sides of use as we have, but *ontologically*. It is thought that observable behavior of a human body is *ontologically* independent, and falls short of, what they are evidence for, i.e., other minds. This is also how Bird perceives the situation with laws of nature: the instances observed simply fall short, again, ontologically, of what they are meant to be evidence of: the *law-like* connections. Bird emphasizes this “ontological gap”, not as a motivation for employing the criterial strategy, but as a condition of adequacy for any account of laws of nature, on the ground that only with an ontological gap can it become intelligible how laws of nature can *explain*, rather

than merely subsume, its observed instances. Despite a lack of explicit recognition on Bird's part, wanting to achieve an ontological gap is precisely what makes him think that criterial strategy can be fruitful.

How much gap is the “ontological gap”? In both the case of other minds and the case of laws of nature, the ontological gap is such that, given the evidence-fact, the fact to be established may still not obtain. There is, in other words, a degree of metaphysical *independence* of the fact to be established from the evidence-fact. But if we recognize this much ontological gap, we have to give up on the *conclusiveness* of the criterion. This is why it is granted, in most attempted applications of the criterial strategy, that criterion need not “confer certainty” on the to-be-established claim, that, in other words, criteria are defeasible. But does the mere admission that criteria are defeasible enable us to “solve” the epistemological problem that prompted the criterial view to begin with: namely the problem that the best evidence falls short of the to-be-established fact ontologically? We could also formulate the question in terms of disharmony. For the ontological gap is just the gap between the verification side of use and the content of a claim that is supposed to license a certain strong consequence, and *that* gap, is just the disharmony of use. So our question has another variant: does the mere addition of defeasibility to the verification side of use – in our case the observation rule – make the disharmony unproblematic?

First note that not just *any* kind of defeasibility will help. For, defeasibility is a feature of *any empirical* warrant for an *empirical* claim. It is in the worldly nature of an empirical claim that there *cannot* be conclusive empirical warrant for it. If my warrant for claiming that there is a red cup on my table is my seeing one, that warrant can be defeated by all sorts of circumstances, ranging from eye problems, to the lighting conditions, to the possibility of a chameleon color-changing cup. If my warrant is being told by someone, that warrant is no less defeasible. But if so, defeasibility *per se* cannot be thought of as a reflection of the ontological gap between warranting circumstances and warranted claim. Rather, it is a quite general feature of the *empiricness* of the claim and warrant.

On the *other* hand, the criterial strategy would lose its promise altogether if the notion of defeasibility is stretched too far. For, *obviously*, no one disputes that, the fact that all the As observed so far are Bs lends *some* support to the law-claim linking As and Bs, at least to the extent that it is *more* supportive of the law claim than if some observed As had turned out to be non-Bs. But that is a very weak sense of

Part II: Semantic Functionalism & The Problem of Laws

“support”. In that sense of defeasible support, the fact that all the gold spheres observed so far have a diameter less than 1 mile lends defeasible “support” to the claim that it is a law of nature that all gold spheres have a diameter less than 1 mile. But that would be a very bizaare way of putting things. For, which circumstance has “defeated” that support in the case at hand? Our intuition about the case is that observations of gold spheres are *too weak* as evidence for the (false) law-claim, rather than that the observational evidence *happens* to be defeated in this particular case. Simply insisting on *calling* such cases of epistemical falling short on the part of observations cases of “defeated” support does not *make* it so.

Put it another way, we should not attempt to make the notion of defeasibility do real work by adopting an *exceptionally easy standard of defeat*, and say that *that* is what distinguishes the cases of ontological gap from the rest of the empirical concepts whose empirical warrants are generally defeasible anyhow. For if we did, there would be no telling how that is different from admitting that the defeasible criterion is just *too weak* for what it is meant to be a criterion for, *even in cases where the claim to be justified is true*. Admitting that, however, is just admitting the epistemological problem that one wants to *avoid* by adopting the criterial view. The situation is not better when looked at from the harmony-of-use perspective. For, saying that the evidence sanctioned by the observation rule is very easy to defeat amounts to saying that the observation rule is *not reliably usable*. So, instead of having disharmony between a reliably usable rule on the verification side and rules on the pragmatic side, admitting defeasibility liberally would put us in the *worse* situation of having no reliably usable rule on the verification side at all!

No one, including Bird, has provided any reason to believe that there is an intermediate notion of defeasibility that does the work of bridging the ontological gap and disharmony: a notion that is *more* than the generic defeasibility applicable across the board to empirical warrants for empirical claims, but is *less* than arbitrarily labeling any epistemic falling short as “criterion defeated”. What we need is a convincing and principled difference between some criterion’s being defeated and that criterion’s being, as such, too weak for the claim for which it is a putative criterion. The criterial view advocates do not provide such a difference. Nor is there any independent reason to think their strategy can be successful. For if nothing is done about the basic facts of evidence and consequence, and so nothing is done about the basic facts of disharmony, why should the addition of defeasibility bring the same

two sides any closer to each other? In a way, Bird realizes the futility of the criterial view when he admits that he does not have a solution of the induction problem.¹⁸⁵ But his evident insouciance does not catch up with the grimness of reality, which is that unless we can see our use of the law-operator as having two harmonious sides, we cannot even be certain of having a *coherent concept* of laws of nature.¹⁸⁶

§4.3.2.3 Realism Backed by Analysis and Truth-Conditions

I introduced the criterial view by saying that it made the right first step by focusing on *use* rather than on truth-conditions. The strategy fails because it does not effect a fundamental change of the disharmonies that plague the concept of laws. The majority of nomological realists, however, present their accounts in terms of an *analysis* of what laws of nature *are*, or, in terms of truth-conditions for law statements. In this section, I review briefly two prominent attempts to do that, and especially the critiques waged by an anti-realist, van Fraassen. Though I fundamentally disagree with their outlook on laws of nature, it is often the anti-realists who most sharply perceive the disharmonies of use and the insurmountable difficulties they present to a realist conception of laws. By this brief review I hope to render vivid how *analyzing* the concept of law in terms of non-nomological items is not only powerless to remove the irritation created by the basic facts of disharmony of use, it also tends to lead to distracting metaphysical disputes.

§4.3.2.3.1 Necessitarian Analysis

According to one kind of analysis, the so-called “necessitarian” view, a law of nature is a kind of necessary truth, where a necessary truth is a truth that holds at all possible worlds that are “accessible” to the possible world that is actual. There may be other refinements, for instance additional requirement of universality. But the basic idea of this approach is to analyse laws in terms of a special, “nomic” sort of necessity, and then analyze the latter in terms of a special, “nomic” accessibility relations among possible worlds.

¹⁸⁵ See the passage of Bird’s *Laws and Criteria* on page 531, quoted above in the section titled “Realists”.

¹⁸⁶ The McDowellian response to the criterialists for other minds is hardly available for us. For it is hardly a plausible thing to say that we can *directly perceive* laws of nature. The correct response here, as I shall argue below, is to correct our conception of what counts evidence for and consequence of law statements.

Part II: Semantic Functionalism & The Problem of Laws

The obvious motivation for this kind of analysis is to *secure the modal/counterfactual rules* for the law-operator. It *follows* straightforwardly from such an account that law statements license certain counterfactual and modal statements, which agrees with our practice. However, necessitarian accounts have faced many objections. One of them, for example, is that they cannot account for vacuous laws (laws that do not have any instances in the actual world) any better than the anti-realist position.¹⁸⁷ One other popular objection, perhaps less focused on periphery-phenomena concerning laws, comes from van Fraassen, who argues that the necessitarians, like the universalists to be discussed below, suffer from the so-called *identification problem*. The identification problem is the problem of saying *which* of many formally admissible accessible relation is *the right one*, the one that serves as the basis for laws of nature. As van Fraassen formulates the argument:

Certain characteristics of that relation [i.e. the nomic accessibility relation] may be postulated, for example that it is reflexive. If we make the list of postulates long enough, will that single out a unique relation? No, it won't, unless it is one of those trivial relations which either hold between all worlds or between none. Otherwise we can always find a distinct, isomorphic relation, which satisfies the same postulates. ... We can't single out the relation by description; and obviously we also can't by pointing to it. This is the *identification problem*.

– van Fraassen, *Laws and Symmetry*, page 72.

Following the identification-challenge cited here, we can ask the necessitarianist to choose from a number of alternatives for identifying the nomic accessibility relation: deictically, by using a proper name (this one van Fraassen did not mention, but let us throw it in the package for our necessitarianist), or by description. It is quite obvious that deixis and proper name are pseudo-options. So the only option open to the necessitarianist is to come up with some sort of description. Van Fraassen concludes, apparently from certain formal characters of possible worlds, that no description can uniquely fit one accessibility relation.

Now it is not immediately clear why van Fraassen thinks this is a problem. On the one hand, if true, the conclusion of the argument merely reflects on the limitation of the descriptive resource of our language, not on the plausibility of the

¹⁸⁷ On this point, see, for example, D. H. Mellor's *Necessities and Universals in Natural Laws*.

necessitarianist analysis of laws of nature. On the other hand, at least on one reading, the conclusion is obviously false. For, we can describe “the nomic accessibility relation” this way: the accessibility relation with the property that, a possible world w_1 is accessible by this relation to another possible world w , just in case for any law of nature L which holds in w , its relevant instances in w_1 are true. In fact, the phrase “the nomic accessibility relation” might be thought as more or less an abbreviation of some such definite description. Van Fraassen might, if this is the answer to his identification-challenge, complain that the necessitarianist isn’t really giving an “analysis” of the notion of laws of nature, if the only way the necessitarianist can tell us which accessibility relation he means is by employing blatantly nomic-vocabulary. But why *should* the necessitarianist be able to say, with words not itself already nomic, which accessibility relation underlies laws of nature, as opposed to some other (say deontic) modal facts?

The real worry expressed by the challenge “tell me *which* accessibility relation it is that underlies laws of nature, without using nomic vocabulary” appears not to be one of *explanatory* circularity, but an *epistemological* anxiety. The challenge is essentially a rhetorical question: how does the introduction of nomic accessibility relations among possible worlds by the necessitarianist help us to *know* which accessibility relation it is? If we did not have a solid epistemological foundations for our laws-of-nature claims, how does the introduction of something *otherworldly* help us bridging the gap?

On this reading of Van Fraassen’s identification problem - the only sensible reading in my view - therefore, the complaint is directed at the epistemological side of the disharmony of use. If laws of nature fall into disrepute because our purported descriptions of them have an inconsistent use-pattern, then analyzing laws in terms of relations between possible worlds only *dramatizes* the problem by making the inconsistency of use even more obvious: if it is hard enough to base claims with predictive power on singular observational claims about the past, how much harder it must be to base claims about *other possible worlds* on claims about the *actual world*? For, if laws of nature claims can be supported by observed non-modal instances in the actual world, and laws of nature consists in what is the case at other independently existing possible worlds related to our world via some primitive accessibility relation, then it must be possible that some counterfactuals (what is the case in other possible worlds) can be supported by non-modal facts in this world. This is just the

Part II: Semantic Functionalism & The Problem of Laws

disharmony between non-modal evidence and counterfactual consequence of a law of nature claim. The force of Van Fraassen's rhetorical question about identification, on this reading, is to make the epistemic remoteness of other possible worlds evident.

§4.3.2.3.2 *Universal Analysis*

A similar problem plagues the so-called universal-accounts of laws of nature. According to such analyses, whose prominent defenders include Dretske, Tooley, and Armstrong, a law of nature is a *relation between two (or more) first-order universals*. Obviously, it cannot be just *any* relations. The relation has to be, on the one hand, such that if it holds between A and B, then all instances of A are instances of B. On the other hand, if the account is to achieve a *distinctive* kind of law contents, the relation cannot be *just* that: it cannot be what van Fraassen calls the relation of *extensional inclusion*. In fact, precisely because one of the key motivations of this type of analysis is, like the necessitarian analysis, to secure the modal-consequences of law statements, the relation in question has been christened “*nomic necessitation relation*”.

Van Fraassen asks of this type of analysis the same identification-question: *which* relation between two universals is the one required? He emphasizes, rightly, the fact that the universal A and its instances are two *different kinds of things*, so that an “intimate” relation between two universals in no way guarantees a close relation between their instances. In particular, the requirement that all As are Bs is not guaranteed to be met by just *any* relation, however close, between the *universals* A and B.¹⁸⁸ There are attempts to identify this nomic necessitation relation in a way that satisfies this basic requirement. So Tooley gives the following descriptions that collectively are supposed to uniquely fix that relation (called “N” here):

- (i) N is a *binary* relation among universals;
- (ii) N is *irreducibly of order 2*, relating universals of order 1 (universals of order 1 have particulars as instances);
- (iii) N is a *contingent* relation among universals;
- (iv) If N holds of A and B, that fact *logically entails that all As are Bs*.

¹⁸⁸ The difficulty of meeting this requirement is what van Fraassen calls the “inference problem”.

Van Fraassen is not satisfied with these. He thinks there is a regress problem. Essentially, he thinks that no relation relating A and B, irreducibly of order 2, can *logically* entail that As are Bs. So he reckons that a yet *higher order* relation is needed to ensure this non-logical entailment, something along the following lines: The relation N* of order 3 holds among the three universals N, A, and B, just in case, if N holds of A and B, then it follows that all As are Bs. But to fix this relation N*, which again is not a logical one, one needs another relation N** of order 4 and so on *ad infinitum*.

Here van Fraassen seems to have missed the intent of the proposal, even though his worry does have understandable causes. The “logical” entailment requirement (iv) is meant to say that it is part of the *logic – or meaning – of the very term “N”* that its holding between A and B should entail that all As are Bs. The role of (iv) is somewhat like the role of an introduction-rule in the fixation of the meaning of a logical constant, except that the meaning to be fixed by (iv) – together with the rest of the descriptions – is not that of a *logical constant*, but that of the nomic necessitation expression “N”. Understood this way, the proposal still has problems. One of them is that, as Alexander Bird correctly points out,¹⁸⁹ there is no reason to think that (i) – (iv) *uniquely* fixes a relation.

But the real problem is again the recalcitrant disharmony that refuses to go away. Amongst these conditions (i) – (iv) the only condition that makes the proposal a *nomological realist* one is condition (ii). It ensures that the N is not merely the extensional inclusion relation, and the secret hope is that (ii) will help select out those that have *modal* consequences. But what sort of *evidence* are we going to get for this *irreducibly second order* relation? First order facts, facts about particulars, seem to be completely inadequate for establishing anything that is *irreducibly* higher order. For if they did, would not the fact established be *reducible* to a fact of the same order as the evidence-facts? On the other hand, it seems clear that we cannot directly *observe* facts of a higher order: we can observe that an A is an B, but not that a 2-order relation holds between the *universal* A and the *universal* B. What we have then, is a manifestation of the disharmony having exactly the same structure as we have just seen with the necessitarians:

¹⁸⁹ C.f. Bird’s *Laws and Criteria*, 532pp.

Part II: Semantic Functionalism & The Problem of Laws

- (a) we adopt an analysis of law statements to give them a kind of *contents* so that the statements are guaranteed to satisfy the modal/counterfactual rules of use;
- (b) the adoption of such contents then leads to an acute *epistemological problem*, for it becomes clear that no observational evidence can ever give us warrant for such contents.

The case of necessitarians is a bit worse, for the epistemological problem it helps generate also takes the form of an *ontological* inscrutibility of one of the explanatory primitives the account introduces, namely, possible worlds. But the key reason for failure is that neither type of accounts tries to remove the underlying disharmony of use, and so inevitably fall into the steps (a) and (b).

§4.3.2.3.3 Why Realism-by-Analysis Fails

The failure can be illustrated by an analogy. We have seen that the disharmony of the use associated with Dummett's example "Boche" makes the expression incapable of expressing a coherent concept. Suppose now that someone defends the concept of a "Boche" by giving it an *analysis*, along the following lines:

Forget the fact that the use-rules for "Boche" licences the inference from being German to being prone to cruelty. If you find the inference suspicious and question the availability of such a concept on that ground, I can tell you *directly* what a "Boche" is. A "Boche" is just a person with a *teutonic* character.

How convincing is this "direct" analysis of the concept of a "Boche" for those of us who perceive a disharmony of its use-rules? The analysis assumes an unexplained explainer, a certain "teutonic character", and we have *no idea* what that is. We might, with van Fraassen, say that the analysis faces an "identification problem" of specifying what a "teutonic character" is. But the problem comes down to this. *Whatever* it is, if the "teutonic character" is enough to ensure tendency to cruelty, then we cannot possibly establish that someone has a teutonic character by merely noting that she is German. That this is the case is of course pre-programmed from the beginning, by the facts of disharmony.

The right way for a realist to approach the debate with anti-realist about laws of nature is not, therefore, to attempt a *direct analysis* of what laws are, in terms that

are not explicitly nomological. The tendency for skepticism about laws of nature to flare up every now and then has root in the fact that the picture of language use concerning law statements presupposed by us *all* contains very disturbing disharmonies that, if left untreated, will *ensure* an abundant supply of anti-realists. The very *first* step for a successful defense of nomological realism must be to remove this perpetual temptation to anti-realism, which means, to remove the perception of disharmony. If that is not possible, than realism about laws of nature is doomed. On the other hand, if nomological realism *is* correct, than it *must* be that the picture of language use I have sketched at the beginning of my brief survey is inaccurate. It simply cannot be that the expression for laws of nature has a set of principles of use that contains disharmonies. To sketch the *correct* picture of language use for the law-operator is the task of the upcoming chapter.

Part III: Laws and Skills – A Therapy

§5 Elements of Nomological Practice: Experiments, Practical Deliberations, and Skills

§5.1 Introduction	194
§5.1.1 Correcting the Traditional Picture	194
§5.1.2 Binary Law Statements.....	196
§5.1.3 Schematic Regimentation For Law Statements	197
§5.2 Desiderative Inferences and the Practical Deliberation Principle.....	199
§5.2.1 Varieties of Inferences	199
§5.2.1.1 Normative Characterization	200
§5.2.1.2 Semantic Characterization	202
§5.2.1.3 Taking Stock	205
§5.2.2 The Practical Deliberation Principle.....	206
§5.2.2.1 Permissive vs. Commissive License	207
§5.2.2.2 From Counterfactuals to Desiderative Inferences	209
§5.2.2.2.1 Advantages of Desiderative Inference over Counterfactuals	209
§5.2.2.2.2 Methodological Equivalence	211
§5.2.2.3 Contextual Parameters & Implicit Understanding.....	215
§5.3 The Skill-by-Experiment Principle	218
§5.3.1 Skill by Experiment	218
§5.3.2 Skill-by-Experiment As Warrant.....	220
§5.3.2.1 The Phenomenology of Experimental Success	222
§5.3.2.2 Motoric Capacities & Spontaneity	223
§5.3.2.3 General Subject-Matter & Repeatability	223
§5.3.2.4 Generality of Warrant.....	227
§5.3.3 Some Contrasts	228
§5.3.3.1 Contrasting With Perceptual Warrant: Public Accessibility	228
§5.3.3.2 Contrasting With Universally Quantified Claims: Resilience & Non-Additivity	230
§5.3.3.3 The Inversion of Normative Standings (& The Light that Dawns Over the Whole)	232
§5.4 The Problems of Disharmony Revisited.....	236
§5.4.1 The Problem of Induction.....	236
§5.4.2 Presumption of Uniqueness of Meaning	237
§5.4.3 The Counterfactual Problem	238
§5.4.4 Implicit Harmony.....	238

§5.1 Introduction

§5.1.1 Correcting the Traditional Picture

Statements of laws of nature express a distinct kind of contents, and some of them are true. In order to see how that can be the case, one need first of all to see the use we *actually* make of law-statements is different from the picture described in §4, which is a picture that contains conspicuous disharmonies. If we call that picture the **Traditional Picture**, the first step towards achieving our goal is coming to see that the Traditional Picture is in fact inaccurate.

The main step in doing that is a defense of the central thesis of this chapter, which is that the core of the language use concerning law statements consists of the following two basic components (these will be expanded later, but they do form the essential core):¹⁹⁰

Verification Side: Skill-by-Experiment Principle:

The warrant for a law of nature statement can be secured by mastering a certain kind of complex skill through experiment;

Pragmatic Side: Practical Deliberation Principle:

A law of nature statement licenses¹⁹¹ the making of certain kind of steps (to be called *desiderative inferences*) in the course of a practical deliberation.

Detailed formulations of these principles require lengthy discussions about skills, experiments, and what I shall call desiderative inferences, to which we shall come presently. Before doing that, let me give a general idea of what this switch of gear means.

First of all, neither principle relates law statements with other *statements*. Steps in a practical deliberation or the mastering of a complex skill are not statements, though both can involve some sort of linguistic capacities. Consequently, neither of the two principles is strictly speaking an *inferential* principle, the way both the introduction and elimination rules for a logical constant, as well as the observation

¹⁹⁰ For discussions of the Dummettian terminology “verification” and “pragmatic” side of use, c.f. §4.1.1 above.

¹⁹¹ For clarification for the *kind* of license meant here, see discussions on permissive vs. committive licenses below.

rule and the counterfactual rule for law statements we discussed earlier are truly inferential principles. This is one of the most profound differences between the picture of language use I propose here on the one hand and the Traditional Picture on the other. I will have more to say about the significance of this feature of change in §7. For now, the following remarks should suffice. On the verification side, the warrant obtained through the acquisition of skills by experiments is not a matter of justification by gathering *evidence*. The *type* of warrant here is closer to, but essentially distinct from, the type of warrant a *perceptual experience* provides a subject with for an observation claim. The point is not just that experiments are not statements, but rather also that, more broadly speaking, neither the status of having acquired a skill nor the process of acquiring it by experiments has *representational purport*.¹⁹² So the relation between that which provides warrant and that which receives warrant is not between claims, so cannot be inferential. Similarly though, the *consequence* licensed by the practical deliberation rule is not a statement, nor a representational act of some other sort, either. So the relation between a law statement and its consequence sanctioned by the practical deliberation rule is likewise not an *inferential* one. This is why the realist solution I offer based on my diagnosis in this chapter will not be an inferentialist one in the narrow sense.

Second, to say that the two principles mentioned form the essential core of the linguistic use we make of the law-operator is not to completely jettison elements of the Traditional Picture. For instance, the deduction rule – the rule that we can justify a law of nature claim by deductively deriving it from another law of nature claim – did not figure in either of the two disharmonies we identified in the Traditional Picture, and so it may stand or be qualified or fall on completely different grounds, and need not be affected by my proposal at all. Nor does the explanation rule suffer a purge. Elements in the Traditional Picture on the pragmatic side – the modal and counterfactual rules – are as we shall see fine in themselves. They are replaced for certain methodological reasons. Even elements on the verification side – the observation rule – may stand, though no longer as a rule on the *verification* side, after the following modification: it is not that observed instances have no role at all in our

¹⁹² Though as we shall see, experimental success can contain “information” without purporting to represent. See the discussion on explicit making in §7, especially §7.2.4 below.

Part III: Laws and Skills – A Therapy

search for and *formulation* of laws of nature, but they are not an adequate source of *warrant* for law statements, nor the type of warrant characteristic of them.

The extent of “correction” over the Traditional Picture my thesis represents is thus not as far as it seems. It does however completely replace the old observation rule on the verification side. For acquiring skills by experimentation is a genuinely different kind of warrant from observation of particulars. By contrast, I am of the view that it is *essentially correct* to view counterfactuals as the distinctive consequences licensed by law statements. This is because counterfactuals, of the type intended in the Traditional Picture, are roughly equal in strength as their replacement I propose, desiderative inferences. But I choose desiderative inferences instead of counterfactuals in my formulation of the pragmatic side of use because, first, desiderative inferences do not suffer from the technical difficulties that plague counterfactuals – messy semantic analysis and disambiguation needed to identify the *right kind* of counterfactual contents intended – and second, more importantly, desiderative inferences make the “matching” with the verification side more transparent than counterfactuals or modal statements would.

The rest of this chapter is devoted to detailed exposition and defense of these two principles of use, beginning with the practical deliberation principle. At the end of the chapter, it will emerge straightforwardly that there is in fact *no* disharmony in the way we use law statements.

§5.1.2 Binary Law Statements

I have been using the schema ‘L(G(A, B))’ for a typical law statement, where ‘L’ stands for a law-operator – say ‘It is a law of nature that’ – while ‘G’ forms a general claim out of two general terms ‘A’ and ‘B’, perhaps in the form ‘All As are Bs’. Under this schematic representation, the central problem we diagnosed in the previous chapter and are now prepared to fix is the disharmony in the principles of use that seem to govern the law-operator ‘L’. This regimentation of law statements is very crude for sure. But they are not more crude than the symbolisms widely used in the literature on the realism-vs.-anti-realism debate on laws of nature. That is hardly something to wonder about. The most fundamental issue is whether and how we can come to see *any* expressive device as intelligibly expressing a universal content genuinely stronger than that of universally quantified claims. The more fine-grained

syntactic structures of *real life law statements* need not unnecessarily complicate *this* task.

Although more details of form will be taken into account in the positive portion of my story (§7 below), I shall restrict myself from now until the end of Part III to the simple form of law statements having a *binary* form:

It is a law of nature that, any A is also a B.

The regimentation of binary law statements will be further refined below, but it will retain the binary form. The way to think about binary law statement is *not* to think of them as *abstractions* or *simplifications* of actual law statements. I intend to think of them as an early stage and a component of the rational reconstruction of actual nomological practices. Though this nomological practice does not yet have the sophistication of modern science – to which we shall come in §7 – it is sufficient for demonstrating the illusion of disharmony that irritates realists and anti-realists alike.

The contrast between this simple nomological practice and the full sophistication of science can be formulated in terms of the following properties of binary law statements:

- a) strongly qualified by *ceteris paribus* clauses;
- b) simple: has only two terms, ‘A’ and ‘B’;
- c) isolated, rather than belonging to a system, inter-connected through, say “derivation” relations;
- d) no essential use of mathematics, nor other kinds of numerical device.

We will come back to these features and their counterparts for scientific laws later in §7. Until then, all “law statements” are binary law statements unless otherwise stated.

§5.1.3 Schematic Regimentation For Law Statements

To refine the regimentation of binary law statements, take our previous example about copper and electricity. The presence of the dispositional predicate “conducts electricity” means that the law statement can be unpacked into:

It is a law of nature that, when a piece of copper is subject to electric potential differential at a time T, then it will carry electric current at time T.

Part III: Laws and Skills – A Therapy

Intuitively, we can think of this statement as talking about a *type of situation S*. Our schematic formulation does not depend on a precise definition of situations or situation types. But we should think of a situation type as a particular kind of constellation of objects bearing certain properties and relations to each other. In the case at hand, a situation of type S is one in which we have *a piece of copper existing at a time T*. This situation type provides for two objects, one of them bearing the property of being copper, the other bearing the property of being a moment in time. The law statement says *of this type of situations that: it is a law of nature that, any situation of type S that has the further property A also has the further property B*. A and B are two additional properties that can be born by a situation of type S. The two additional properties linked up by our example law statement are: the copper piece being subject to an electric potential differential at the given time, and the copper piece carrying electric current at the given time.

To generalize and formulate the thought schematically, we can write the following for a law of nature statement:

It is a law of nature that, given S(x ... z), if P(x...z), then Q(x...z).

Here the lower case letters ‘x’ etc. stand for objects. ‘S’ is a description schema that says of the objects designated by ‘x’, ... ‘z’ that they form a situation of type S. ‘P’ and ‘Q’ designate the two additional collective properties the objects that form the situation of type S can have. For our copper-conducts-electricity example, these various predicates are:

S(x, y): x is a piece of copper and y is a moment in time;

P(x, y): x is subject to electric potential differential at time y;

Q(x, y): x carries electric current at time y.

The law-statement schema can either be symbolized further, say in the form of (‘LS-S’ stands for ‘Law Statement – Symbolic form’):

(LS-S) $L(\forall x \dots \forall z [S(x, \dots z) \rightarrow \{P(x, \dots z) \rightarrow Q(x, \dots, z)\}])$

or it can be further abridged with minimal symbolism, as in:

(LS) *It is a law of nature that, if any S-situation is A, then it is also B.*

Compared with the traditional regimentations of law statements – in terms of a law-operator combined with two general terms – mine has an extra element of complexity, namely the “situation of type S”. The reason for this additional complication is not to track surface grammatical features of actual law statements *per se*. It is that this complexity is necessary for the kind of analysis of linguistic use associated with law statements that I believe is adequate.

In the following, I will first explain the basic elements that make up the core of our linguistic practice *in themselves*, before linking them up with the law statements. Why the schematic regimentation proposed here makes sense will be clear in the second of those two steps.

§5.2 Desiderative Inferences and the Practical Deliberation Principle

§5.2.1 Varieties of Inferences

Desiderative inferences are a kind of inferences made in the course of a practical deliberation leading up to action. The most general form of the kind of such inferences that is of the most relevance to us is:

I want that X; so, I shall see to it that Y;

The underlined verbal parts are crucial to – though as we shall see presently, not sufficient to guarantee – the inference’s being a different kind from the most widely discussed sort occurring in theoretical reasoning. What is traditionally called an inference is a reasoning process from premises to conclusions in which the author of the inference aims at deciding *what to believe*. For that reason, I shall call the traditional sort of inference *doxastic* inference. By contrast, through an inference of the kind represented by the schema above an agent aims to establish *what to do*. It would be somewhat misleading to call such inferences “conative” inferences, for the family of words based on the word-stem “conat-” have come to be associated with aspects of the psyché such as impulses and desires so that ‘conative *inference*’ would have appeared to be an oxymoron: a *discursive* action that is at the same time an

impulse. For this reason I will coin a new term for such inferences: *desiderative inferences*.¹⁹³

§5.2.1.1 Normative Characterization

I have been characterizing the difference between doxastic inferences from desiderative inferences in terms of what happens in the corresponding *psychological* process: in one case a new *belief* is formed, in the other case, a new *intention* or *desire*. But the difference is also clear in the normative dimension of justification. Consider the following pair of inferences:

- (1) I want to crack open the nut; so, I shall hammer it;
- (2) John wants to crack open the nut; so, he shall hammer it;

Superficially, the two inferences look quite similar. They differ only in what occupies the subject position of their respective clauses. Let us suppose furthermore that the inference is made by John himself, so that the *referent* of ‘I’ in (1) is exactly the same as ‘John’ in (2). The similarity is an optic illusion.¹⁹⁴ Let us look at them in turn, beginning with (2). The most plausible reading of (2) is that the speaker infers from John’s wish to open the nut to a prediction of his behavior: hammering the nut. To make inference (2) is to predict what someone will do based on what he wants. How might one *support* such an inference? One way of supporting such an inference is by

¹⁹³ It is important to emphasize that the new category being discussed here is not one of *statements* or *sentences*, notwithstanding close relationships between a specific kind of statements - namely conditionals - and inferences. There are some strategic reasons for choosing inferences over conditional statements to characterize the consequence of application for law-of-nature statements. One of them is the clarity of harmony that results, between practical skills and practical inferences. Some of the other advantages are discussed in section §5.2.2.2.1 below.

Still, there is in some sense an equivalence between some classes of conditionals and corresponding classes of inferences (for example, c.f. Brandom’s discussion in *Making It Explicit*). In section §5.2.2.2.2 *Methodological Equivalence* below, I discuss the sense of equivalence between desiderative inferences and counterfactual conditionals. Another class of conditionals in the vicinity are the so-called “anankastic conditionals”, conditionals of the form “If you want X, you have to do Y”. It might be thought that these correspond to desiderative inferences of the form “I want X, so I shall see to it that Y”. The sense in which this is true is that a causal relation between Y and X underlines - in some generic sense of “underline” - both. However, the anankastic conditional is true only when doing Y is a *necessary condition* for X, which is not required for the validity of the corresponding desiderative inference.

¹⁹⁴ I say “optic” here because the illusion arises in large part due to the way in which (1) and (2) *look* similar to each other when written down. But strictly speaking, the similarity itself is not optic; it is structural. What is optic is the particular mode of presenting the similarity in the medium of writing. In any case, misleading structural similarities in language such as these are key examples of what Wittgenstein calls the “bedevilment” of the understanding through language.

(a) pointing to the fact that John believes hammering is a good way of cracking open a nut, and (b) noting that given this belief on John's part and his wish to crack open the nut, it follows from *certain general psychological truths* that John is likely to hammer the nut.

Now consider the inference (1). In one sense of "semantic content", the premises of (1) and (2) have *exactly the same* semantic content, and so do the conclusions of (1) and (2). But the strange thing is, (1) still is not equivalent to (2). It is true that (1) has *one* reading that is exactly the same as (2), namely the *prediction-based-on-psychology* reading, or, for short, the prediction reading. According to the prediction reading, John, in making the inference (1), is predicting his *own* behavior based on one of his wishes. One *can* understand the inference (1) this way, but it is seldom the case that a person would try to gain *self-knowledge* about his intention in this manner: by first introspectively examining what wishes he has, and then infer from them a prediction about what he himself will want to do.¹⁹⁵ Clearly, the prediction reading is not the most natural reading for (1). The most natural way to read inference (1) is to understand it as the verbal expression of John's *process of making up his mind about what to do*, that is, to understand it as part of a *practical deliberation*. Let us call this reading the *practical deliberation reading*. Now, under the practical deliberation reading, what would be an appropriate justification for inference (1)? First of all, the right form of *challenge* to (1) under this reading is not 'no, you are mistaken; that will not happen'. One way of challenging it is to say: no, you *shouldn't* do that, hammering the nut won't get you anywhere.¹⁹⁶ To *support* the inference (1), John might say something like "well, hammering *is* a good way to crack open the nut". Note that, in contrast to the prediction-reading, John will *not* support his inference by citing the fact that he *believes* that hammering is a good way to crack open a nut, nor any general psychological truths about agents with desires and beliefs. What John needs in order to defend the inference under the practical deliberation

¹⁹⁵ If it is hard to read the inference (1) this way – and it *is* hard indeed – consider a similar inference "I want to buy a BMW; so, I want to buy a car". Here the occurrences of 'I want' are merely descriptive. No new desire is formed. Rather, the inference is an *analysis* of an existing desire.

¹⁹⁶ Anscombe, in her classic discussion on the subject of intention, says that the appropriate challenge to a declaration of intention is a counter-declaration to prevent the subject from succeeding. The kind of challenge that I am considering now is not a *practical* challenge in Anscombe's sense. Instead, the kind of challenge I am interested in is the kind to which one can appropriately respond by trying to *justify* the goodness of the *move* from a desire to the intention. For details of Anscombe's discussion, see her *Intention*.

reading is a *fact about hammering and nuts*, not facts, particular or general, about psychology.

§5.2.1.2 *Semantic Characterization*

Pragmatics & Doxastic Bias

In addition to the psychological and normative ways of drawing the difference between (1) (under the practical deliberation reading) and (2), there is a third, broadly *semantic* way. To explain it, first consider what I call the *doxastic bias* of the declarative form of an utterance. It has often been remarked that the following two statements are equivalent:

- (3) John wants to crack open the nut
- (4) *I believe that* John wants to crack open the nut.

This equivalence has been taken up by philosophers as a puzzle: how can two statements, one of which contains the other plus some *extra words*, be equivalent? The solution to the puzzle is usually that the equivalence between (3) and (4) is a matter of pragmatics, and there is a difference between the two in semantic content, strictly so-called. The Frege-Geach imbedding test is then used to show the difference in semantic content.¹⁹⁷ What interests us in the present context is not the *how*-question, but the *why*-question: why is there a pragmatic equivalence between (3) and (4), which are different in semantic content? The answer, I suggest, is that the declarative form of the utterance (3) has a *doxastic bias* in this sense: for someone to utter *p* in declarative form gives us, simply on the basis of this utterance and its declarative form, prima facie entitlement to infer that the person *believes* that *p*, though the same utterance does *not* allow us to infer that the person *wants* or *intends* to see to it that *p*, or *wishes* that it were *p* etc.¹⁹⁸ That there is a bias towards belief is made evident by the *lack* of equivalence between the following:

- (5) John wants to crack open the nut
- (6) *I want/wish that* John wants to crack open the nut.

¹⁹⁷ C.f. for example, Geach's *Assertion*.

¹⁹⁸ The dominant answer to the *why*-question in the literature for philosophy of language is by reference to the existence of some sort of pragmatic axioms (e.g. "do not say what you do not currently believe") governing speech. I am not interested in developing a detailed answer in this direction. What is important in this context is to develop a sufficient understanding of the case to appreciate a pragmatic *contrast* between doxastic attitudes and desiderative attitudes.

If someone utters (5) in normal declarative tone, we are not thereby entitled to attribute to him the wish expressible by him with (6).

Conveying What is Not Said

Given the doxastic bias of declarative statements, we can describe the relationship between (3) and (4) by borrowing a terminology from Sellars. Sellars claims, as we have seen in our earlier discussions on semantic functionalists' account of nomological/modal statements,¹⁹⁹ that a declarative utterance “conveys”, but does not assert, facts about the speaker's beliefs. Using this terminology, we can say that the phrase ‘I believe that’ in (4) describes or asserts *explicitly* what is merely “conveyed” implicitly by (3) through its declarative form. The relation between the asserting function of the declarative form, the latter's doxastic bias, and the “conveyance” function is this: it is precisely because

- a) the default pragmatic function of declarative discourse is to *state* and *assert*, and
- b) a stating or asserting act *conveys*, in standard situations, belief in what is stated on the part of the speaker, that
- c) the declarative form gives rise to a prima facie entitlement to infer belief.

The negative side of the doxastic bias means that declarative utterances *cannot*, by merely being in that form, convey the speaker's *desiderative* states such as desires and intentions. This alone does not explain the necessity of having ‘I want’ and ‘I shall’ in:

- (1) *I want to crack open the nut; so, I shall hammer it;*

For, even though (2), or any other inference composed of declarative statements, is capable of *conveying* desiderative states, the statements in (2) do appear to *say* the same as the statements of (1), as long as ‘I’ and ‘John’ are co-referential. If statements in (1) are descriptive devices that *explicitly* state the speaker's desiderative states, surely statements in (2) do that as well.

Saying What Cannot be Conveyed: Attitudes *De Se*

¹⁹⁹ C.f. §3.3.3, especially the sections on Sellars.

Part III: Laws and Skills – A Therapy

The speaker's desiderative states, unlike their doxastic counterparts, cannot be conveyed, so they must be 'said' or stated. However, not *any* descriptive device will do. Simply referring to oneself will not guarantee the expression of one's *own* desiderative state. This has been observed since at least half a century ago, and is known under the heading of attitudes *de se*, as opposed to attitudes *de dicto*.²⁰⁰ A person may have lost memory of his own name and via the observation by a third person come to the knowledge that he can express by "John wants to crack open this nut". But this is not the knowledge that he *could* express by "I want to crack open this nut".

In summary, a semantic description of the desiderative inferences vs. doxastic inferences is this: a desiderative inference is an inference from a statement of the speaker's desiderative state *de se* such as a desire to another statement of the speaker's desiderative state *de se*.

An Asymmetry: Doxastic vs. Desiderative Inferences

To wrap up our general semantic discussion of doxastic vs. desiderative inferences, we might begin by saying that there is a broadly semantic difference between the expressions 'I want/believe' and 'John wants/believes', even though they can, due to a possible *co-reference* of 'I' and 'John', have exactly the same *meaning*, in some narrow sense of "meaning". The difference is that the third-person variant is *merely descriptive* - in this narrow sense of meaning - whereas the first-person variant is also *expressive* of the speaker's psychological states.^{201 202} This expressive function²⁰³ is redundant when the psychological state expressed is doxastic, due to the doxastic bias

²⁰⁰ For a locus classicus, c.f., the paper by David Lewis [1979], as well as the well-documented references therein.

²⁰¹ Some, such as the adherents of a Fregean style of philosophy, would say that the utterances of "I" and "John", if co-referential, have different *modes of presentation* for the same *reference*. I am not concerned with developing exact and comprehensive semantic distinctions here. Rather, my goal here is using known concepts, which may stem from heterogeneous sources, to characterize a new distinction.

²⁰² The underlying semantic facts for this expressive function - or the power to express attitudes *de se* - may be related to the difference in force of expressions in *explicit performatives* such as "I (hereby) promise..." vs. "John (hereby) promises...", even when "I" and "John" are co-referential. As explicit performatives are not directly related to the project in this essay, I shall not pursue this issue any further.

²⁰³ Some would, such as students of Kaplan, prefer to speak of a dimension of the meaning of "I", properly so-called, and suggest that this dimension of meaning is to be best modeled by "characters". Again, I am not interested in developing a model-theory of semantics here.

of the declarative form. But it is in the current state of English indispensable when the states expressed are desires or intentions.

But we can easily imagine how the expressive function of ‘I want/shall’ would *also* become redundant, if we had a *mood* whose default psychological function is to express desiderative states such as desires and intentions, the way the declarative form’s default psychological function is to express doxastic states such as beliefs. In that case, we could formulate the desiderative inference (1) *without* the ‘I want’ and ‘I shall’ prefix by employing that mood. So if the mood consists in writing exclamation marks around the infinitive form of a verb, the inference (1) can be rewritten as:

(1)* !crack open! the nut; so, !hammer! it.

There might be a question as to whether the subject-less (1)* or the following form with full *propositional contents* is better:

(1)** I !crack open! the nut; so, I !hammer! it.

Either way, the use of the expressive ‘want’ and ‘shall’ would not be necessary given the availability of such a mood.

But the lack of this “desiderative” mood on the one hand, and the redundancy of the expressive function of “I believe” for doxastic attitudes *de se* on the other hand, mean that, there is an asymmetry between doxastic and desiderative inferences as far as their formulations are concerned. Formulating a desiderative inference requires the *stating* of one’s desiderative attitudes *de se*, using “I want to/I shall...” By contrast, formulating a doxastic inference usually requires the *avoidance of explicit statement* using “I believe p...”, preferring the formulation “p; therefore q” instead.

§5.2.1.3 Taking Stock

To sum up the discussion so far, we can say that a desiderative inference is an inference of the form:

I want that X; so, I shall see to it that Y;

where (i) the speaker forms a new decision/intention/desire based on an existing desire; (ii) the inference is properly justified by mentioning items referred to in ‘X’ and ‘Y’ rather than what the speaker *believes* about these items or other psychological

Part III: Laws and Skills – A Therapy

facts; and (iii) the prefixes ‘I want that’ and ‘I shall see to it that’ are not merely descriptive, they “convey” or express the speaker’s desiderative states. These three clauses are, respectively, the psychological, normative, and semantic characterizations of desiderative inferences.

Before closing, we should note that there are other kinds of inferences besides doxastic and desiderative ones. Consider the following:

- (8) John is caught in traffic-jam; so, I shall postpone the meeting.
- (9) Jane wants to get on the bus too; so, I shall hold the door open for her.

These inferences are not doxastic, because in their conclusions the subject forms a new decision/intention. Nor are they desiderative, because their premises do not express a desiderative state of the speaker, rather, they express his *beliefs* (about John and Jane respectively). The relation between the premise and conclusion in these inferences is one of *reason for an action*: the conclusion expresses an intention to perform an action, the premise expresses the reason for that action. The kind of reason in question is *practical* reason, which makes it an apposite choice of terminology to call inferences such as (8) and (9) *practical inferences*. For the sake of completeness, we can now say that there are three types of inferences, which are *doxastic*, *desiderative*, and *practical*, respectively. The nature of practical inferences is a topic for practical philosophy, and will not detain us for our investigation of the nomological discourse.

§5.2.2 The Practical Deliberation Principle

The pragmatic side of the principle of use I proposed says that accepting a law statement has the consequence that one should also accept a certain form of desiderative inferences. Now that we have met desiderative inferences, we can give a more precise formulation to that principle. Suppose we accept a law statement of the abridged form ‘L (of any situation of type S, if A, then B)’, then we must also accept desiderative inferences of the following form:²⁰⁴

I want *this* S-situation to be B; so, I shall see to it that it is A

²⁰⁴ See qualifications on this “must also accept”, see discussion on permissive license below in §5.2.2.1 Permissive vs. Committive License.

The phrase ‘this S-situation’ stands for a concrete description of a situation of type S, where the variables in the S-schema are substituted for by singular terms of various kinds (proper names, demonstratives, indexicals, etc.). If we take the more symbolic form of laws:

$$L(\forall x \dots \forall z [S(x, \dots z) \rightarrow \{P(x, \dots z) \rightarrow Q(x, \dots, z)\}])$$

The desiderative inferences it licenses have the following form:

I want, of $a, \dots c$, of which S is the case, that Q is also the case;
so, I shall see to it, of $a, \dots c$, that P is the case.

The letters ‘ a ’ ... ‘ c ’ stand for singular terms that can be substituted for the variables ‘ x ’ ... ‘ z ’. To continue our previous copper-conducts-electricity example,²⁰⁵ the formulation takes the following concrete form:

I want, of a and T , where a is a piece of copper and T is a moment in time, that
 a carries electric current at T ;
so, I shall see to it that a is subject to electric potential differential at T .

In case ‘ a ’ and ‘ T ’ are demonstratives ‘that piece of copper’ and ‘now’ respectively, we get the following sample desiderative inference licensed by the copper-conducts-electricity law:

I want that *that piece of copper*₁ carries electric current *now*;
so, I shall see to it that *it*₁ is subject to electric potential differential *now*.

The practical deliberation principle is a better way to capture the commonly accepted notion that law statements “support” counterfactuals, or, what I have called the “counterfactual rule” in my analysis of the Traditional Picture. This I shall show presently.

§5.2.2.1 Permissive vs. Commissive License

But before that, I need to add a small commentary to our discussion of the practical deliberation principle, a commentary on the nature of the license of a desiderative inference. We have already seen that the law statement (LS) does not necessary

²⁰⁵ C.f. §5.1.3 Schematic Regimentation For Law Statements.

Part III: Laws and Skills – A Therapy

provide justification for the desiderative inference (DI), for ethical considerations can easily override the conclusion of the desiderative inference. So it would be literally wrong to say that if one accepts the law, one *must* also accept the desiderative inference. Law claims provide only a *prima facie* license to make corresponding desiderative inferences.

The issue of license does not stop here. To adequately characterize the kind of relation between law statements and desiderative inferences, let me introduce a distinction between *permissive* and *committive* license. This is a terminology borrowed from Brandom, though the underlying concept must be somewhat modified.²⁰⁶ Brandom draws the distinction based on his distinction of commitment and entitlement to a claim. An inferential transition is permissive if it preserves entitlement, and committive if it preserves commitment. Brandom says that the distinction corresponds roughly to the traditional distinction between deductive and inductive reasoning. I do not think this alignment is correct. Nor do I want to say, as Brandom does, that “instrumental inferences corresponding to the principle ‘Who wills the end wills the means,’ like inferences generally, come in two flavors: committive and permissive. Some instrumental inferences ... are also commitment-preserving”.²⁰⁷ Whereas Brandom makes the difference a matter of uniqueness of means, I think all transitions of desiderative states are permissive. The part of the Brandomian distinction I want to preserve is most clearly articulated in a footnote of his :

One can be (taken to be) entitled to claims one is not (taken to be) committed to – these are conclusions one is entitled to draw but has not yet committed oneself to.

– Brandom, *Making It Explicit*, footnote 44, page 675.

The idea can be put somewhat differently, in a form that suits my purpose better. If one is permissively licensed (by a law-statement in our case) to do something (in our case: to form a new desire or intention based on an existing desire), one is *not*

²⁰⁶ For Brandom’s definition of “committive” and “permissive” in terms of commitment and entitlement-preservation, see page 168, *Making It Explicit*; for his identification of deductive and inductive inferences as commitment-preserving and entitlement-preserving inferences, see page 132, *Ibid.*.

²⁰⁷ *Ibid.*, page 237.

compelled to do it. This contrasts with the way in which “p & q” licenses the statement “p”, for example. If I am committed to the claim that p & q, I am, *automatically* as it were, compelled by rationality, also committed to the claim that p. But if I want to get across a ravine, and cutting down a tree is the *only* way to do it,²⁰⁸ I am still not compelled by rationality to form the intention of cutting down that tree. In fact, just the opposite, if the tree in question is sacred to a local community, I may be compelled by (practical) rationality *not* to form that destructive intention. (No one who has watched the film *Avatar* would likely disagree.)

So the first feature of a permissive license is that what it licenses is not done or attributed automatically, that no one is compelled to follow that license by standards of rationality.²⁰⁹ The second feature of permissive licenses is that they provide only *prima facie* justification. A permissive license justifies an action only when no other considerations override that license. In the case of the license provided by a law statement for a desiderative inference, it provides the subject with a justification for the inference only if no ethical considerations or considerations about competing desires on the part of other individuals override the concluding intention of that desiderative inference.

§5.2.2.2 From Counterfactuals to Desiderative Inferences

§5.2.2.2.1 Advantages of Desiderative Inference over Counterfactuals

From the outset – that is, independent of the fruits born at the end of my argument for the harmony of use for law statements – there are some advantages to framing the pragmatic side of use for law statements in terms of desiderative inferences rather than, as is traditionally the case, in terms of counterfactual conditionals. One of such advantages is that the availability and forms of counterfactual markers depend on individual languages. For some languages – Chinese for example – it has been claimed that there are *no* consistent morphological and syntactic markers with which

²⁰⁸ This is Brandom’s example of a “committive” practical inference on page 237, *Making It Explicit*.

²⁰⁹ It is compatible with this to acknowledge that there is *some sense* that rationality does recommend the unique means of an end that one desires to achieve. The kind of rationality in question is that it is rational to *follow through* on what one wants, everything else being equal. So what *compels* about any *particular* transition from one’s wish to an intention flows from this *generic* kind of rationality. Nothing about any particular transition compels. This contrasts with the way in which a deductive transition compels.

Part III: Laws and Skills – A Therapy

counterfactual conditionals can be constructed.²¹⁰ By contrast, all languages must admit practical deliberations and in particular steps leading from a desired end to an intended means, which are just what desiderative inferences are. Another advantage of desiderative inferences has to do with the fact that, picking out counterfactuals as the distinctive kind of consequences of law statements is at least *misleading* because it can give the impression that it is essential for nomological connections to be expressed in a contrary-to-fact manner. Many indicative conditionals about the future also express nomological or causal connections that are stronger than mere temporal sequence. Consider: If you talk loudly on the phone, you will wake up the baby. It is true that the indicative conditionals about the future may or may not convey such connections:

If Tom calls, Jane will divorce him.

Depending on the *context*, this indicative conditional might express an *evidential* relation, or a causal one. So on the one hand, in a context where Tom told us that he will call only when his wife informs him of her intention to divorce him, the conditional expresses the *evidential relation* between Tom's call and Jane's decision, without implying that the former *causes* the latter. On the other hand, if the situation is one in which the implicit addressee of Tom's call mentioned in the conditional is an old mistress of him, then the conditional tells us that a call from Tom to his old mistress will be the last straw, and will bring Jane to her senses.

It is arguable as a matter of empirical linguistic fact that contrary-to-fact formulations preclude the evidential kind of reading of a conditional, and for that reason, counterfactual conditionals are the appropriate kind of conditionals for characterizing the pragmatic side of use for law claims. But it is a purely empirical accident that no nomological-marker exists in English for indicative conditionals about the future. If it *existed*, we could have used such indicative conditionals in place

²¹⁰ For discussions on this topic, see A. H. Bloom, *The Linguistic Shaping of Thought*, T. K. Au, "Counterfactuals: In reply to Alfred Bloom", and Feng & Yi, "What if Chinese had Linguistic Markers for Counterfactual Conditionals?". Although Bloom's conclusions about Chinese speakers' ability for counterfactual *thought* are highly implausible, the *linguistic* thesis – which concerns the *means* Chinese speakers have for such thoughts – namely the thesis that the Chinese language lacks linguistic counterfactual markers appear not to be fundamentally challenged. So Feng & Yi write, concessively, "It should be noted that we are not arguing for logically necessary or sufficient markers of CF [counterfactuals], nor do we think they exist". (page 1282, "What if Chinese had Linguistic Markers for Counterfactual Conditionals?").

of counterfactuals to characterize law statements. As far as the *substance* of the matter is concerned, a nomological or causal connection need not be expressed in contrary-to-fact manner. The adoption of desiderative inferences help us shaking off this misleading association. In fact, given an instance of such inferencing, let us say “I want this S-situation to be also B; so I shall see to it that it is also A”, things can turn out *either way*: the subject can (a) *fail* to bring it about the S-situation at hand is also A; or (b) he can *succeed* in doing that. In case (a), we can express the same nomological connection that is expressed by the speaker’s desiderative inference using the *counterfactual* “If the S-situation had been also A, it would have been also B”. In *both* cases, the same connection can be expressed by the indicative conditional about the future *prior* to the speaker’s desiderative inferencing: If the S-situation *becomes* A, it will also *become* B. Although, as we said, the last conditional is subject to an alternative, evidential reading.

The desiderative inference does not, on the one hand, discriminate between case (a) and (b), so does not misleadingly associate nomological connections with the contrary-to-fact case (a). On the other hand, the same desiderative inference precludes an *evidential* understanding: one cannot justify “I want this piece of copper to carry electric current; so, I shall subject it to electric potential differential” by saying that it would be *evidence* for the copper’s conducting of electricity that it is subject to electric potential differential. One justifies it by pointing to the law about copper conducting electricity. So the use of desiderative inference has the methodological advantage that it does not rely on nor is negatively impacted by *linguistic accidents*, such as those concerning indicative and counterfactual conditionals:

- (I) Indicative conditionals can be interpreted both causally and evidentially;
- (II) Counterfactual forms tend to express nomological/causal relations.

§5.2.2.2.2 Methodological Equivalence

But is there a difference in *substance* between employing desiderative inferences and employing counterfactuals to characterize law statements? To see why the answer is no, consider the following three formulations:

- (DI)_{t₀} I want this S-situation to be also B; so I shall see to it that it is also A;
- (CF)_{t₂} If the S-situation *had* been also A, it *would have* been also B;

Part III: Laws and Skills – A Therapy

(LS) *It is a law of nature that, if an S-situation is A, it is also B.*

Let us suppose that (DI), a desiderative inference, is uttered at time t_0 , while (CF), a counterfactual conditional, is uttered at a later time t_2 . Suppose also that both (DI) and (CF) refer to a point in time t_1 – call it the *reference time* – that is in between. That is, we have three times $t_0 < t_1 < t_2$; the S-situation referred to by (DI) at t_0 and by (CF) at t_2 is not yet A nor B at t_0 , so that the desiderative inference (DI) at t_0 makes sense. Now to begin with, we cannot directly compare (DI) and (CF)'s *contents*, for one of them is not a statement at all, but an inference of a special kind, and inferences, being not candidates of truth, do not express propositional content. What is at issue is whether the following is true:

Methodological Equivalence (of Counterfactuals with Desiderative Inferences, Version 1):

Whenever a counterfactual of the form (CF) is truly utterable, a corresponding desiderative inferencing (DI) is justified, and *vice versa*; where “corresponding” means, *inter alia*, having the same reference time t_1 .

Unfortunately, formulated this way, neither direction of the Methodological Equivalence is true without qualifications. The second half is not true because a counterfactual cannot be truly uttered unless its *counterfactual presupposition* is satisfied. So granted that my desiderative inference (DI) at t_0 is justified, it might turn out that I in fact *do* manage to make the S-situation also A, in which case the presupposition of (CF) uttered at t_2 is violated, so that for speakers knowing this fact, (CF) ceases to be a candidate for truth (depending on philosophical persuasion (CF) is either false or does not possess a truth-value). However, the failure of equivalence in this direction is not a genuine threat. For it is due to the linguistic accident that counterfactuals are assertible only when its antecedent A is believed to not obtain. The fact that the desiderative inference (DI) can be justified independent of the truth or falsity of “this S-situation is A” at the reference time t_1 – in particular, independent of whether the agent *succeeds* in making the S-situation A at t_1 – is a *plus*, for it removes the possible erroneous impression that nomological connections are essentially expressible contrary to fact.

The other direction of the Methodological Equivalence fails for a different, more interesting reason. Suppose the counterfactual (CF) is true, the corresponding

desiderative inference (DI) may not be justified. Consider this. I was hungry, and there was a delicious roast-beef sandwich in the fridge which my office-mate brought from home for lunch. It is true that if I *had* eaten my office mate's sandwich, I *would have* assuaged my hunger. And that counterfactual is true because of a nomological connection between eating roast-beef sandwiches and relieving hunger. However, that nomological connection does not equally support the desiderative inference that I could have made "I want to relieve my hunger; so, I shall eat my office-mate's roast-beef sandwich". It is not just that the end does not justify the means (my hunger does not justify eating someone else's lunch), so that the *intention* to eat my office-mate's roast-beef sandwich remains unjustified. The entire desiderative *inference*, which resolves a desire into the intending of a particular means, is not justified. If the means is not justified, then using the means to achieve an end is also not justified.

In general, there is a strong intuition that, while the truth of a counterfactual is only sensitive to *theoretical* facts including laws, the goodness of a desiderative inference is sensitive to *non-theoretical* facts as well, such as ethical norms, competing desires from fellow human beings, etc. This is how it can come to pass that a counterfactual of the form (CF) is truly utterable, yet, due to these other, ethical and moral constraints, one is not justified in making a corresponding desiderative inference (DI). Quite simply, an efficient means for an end is not always a *legitimate*, or justified, means to achieve that end. In fact, desiderative inferences' affinity to ethical factors might appear a decisive argument *against* employing them to make sense of the semantics of law statements, which seem to be purely theoretical statements.

I think it is possible to acknowledge this intuition, yet still make use of desiderative inferences for our account of the nomological practice. What one has to say is that a desiderative inference can be *prima facie* justified, even when other, say ethical considerations ultimately overrule that *prima facie* status. With this proviso, and taking into consideration about the counterfactual presupposition, we can now formulate a modified version of the Methodological Equivalence, with the two modifications underlined:

Methodological Equivalence (Version 2):

Under the assumption that the counterfactual presupposition of (CF) is fulfilled
 – namely that the S-situation referred to did not become A at reference time t_1 -

Part III: Laws and Skills – A Therapy

the counterfactual (CF) is truly utterable *just when* the corresponding desiderative inferencing (DI) is *prima facie* justified; where “corresponding” means, *inter alia*, having the same reference time t_1 .

This second version of the Methodological Equivalence does seem to be plausible. For, given that the counterfactual presupposition of (CF) obtains, both the truth of (CF) as well as the *prima facie* justifiedness of (DI) are equivalent to the truth of the following causal claim (CC):

(CC) The S-situation’s becoming A at t_1 will lead to its becoming B.

uttered at an earlier time, say t_0 . Intuitively, for it to be *prima facie* justified to want to bring about B by bringing about A (*prima facie* justifiedness of (DI)) just *is* for it to be the case that bringing about A will lead to B (truth of (CC)), which in turn just *is*, given that A did not turn out to be the case, for it to be true that *had* A been the case, so would have been B (truth of (CF)). Moreover, the equivalences are not merely material, for they do not depend on the concrete expressions that substitute for ‘A’, ‘B’, and ‘S’. In fact, the equivalences depend only on the meaning of such items as “want”, “shall see to it that”, subjunctive mood, and “will lead to”.

It is true that in asserting the equivalence of (CC) both with (DI) and with (CF) in the way I did, I am appealing to pre-theoretic intuitions, rather than any particular theory about causality, psychology, or counterfactual conditionals. But it would be inappropriate to object to these intuitive claims on the ground that they are not based on well-worked out theories, for, we are concerned with a *condition of adequacy* for any theory about law statements. The most popular way to formulate it is in terms of the counterfactual rule: any theory must make out law statements as supporting counterfactuals. My goal was to make plausible the claim that, this formulation, as a condition to be satisfied by any account of laws of nature whatsoever, is essentially equivalent to the formulation in terms of desiderative inferences. Since the issue is whether it is legitimate to capture a certain pre-theoretical intuition about law statements in a particular form – in terms of desiderative inferences – it obviously must be settled on that level of pre-theoretical intuitions.

The Methodological Equivalence is corroborated by the fact that universally quantified claims are neither strong enough to support counterfactuals nor strong

enough to license desiderative inferences. For example, the universally quantified statement “All dogs born at sea are cocker spaniel cubs”, as long as it does not express a law of nature, does not support the desiderative inference “I want this pregnant dachshund to give birth to cocker spaniel cubs; so, I shall see to it that it gives birth at sea”.

§5.2.2.3 *Contextual Parameters & Implicit Understanding*

The Methodological Equivalence between counterfactuals and desiderative inferences is corroborated in another way. And it has to do with an important feature of counterfactuals, one that has frequently proved an irking nuisance in philosophy: their context dependence. Desiderative inferences come with a feature of its own – they are accompanied by an implicit understanding not related to their verbal expressions – that is the exact counterpart to counterfactual’s context dependence.

To explain, let me begin with the context dependence of counterfactuals. In his pioneering work on counterfactual conditionals Nelson Goodman pointed out that, if we are to analyze the truth-condition of counterfactuals in terms of additional conditions that, when added to the antecedent, entail the consequent “by law”, then we would have problem specifying *which* conditions are to be so added to the antecedent.²¹¹ In David Lewis’s popular semantics for counterfactuals based on possible worlds, the issue takes the following form: which non-actual possible worlds are *closer* to the actual world than others? Somewhat loosely speaking, but hewing close to the core intuition, the problem is this. In evaluating a counterfactual conditional for truth by considering the counterfactual scenarios described by the antecedent, we cannot count on *the conditional sentence itself* to specify which facts must remain fixed, and which facts can go. This “problem” of the counterfactuals is illustrated vividly by Quine’s example of the contrasting pair “If Verdi and Bizet had been compatriots, Verdi would have been a Frenchman” and “If Verdi and Bizet had been compatriots, Bizet would have been an Italian”.²¹² In the first of these conditionals, the fact that Verdi was an Italian is allowed to go, while the fact Bizet

²¹¹ Goodman’s term for this is “the problem of relevant conditions”. See his *The Problem of Counterfactual Conditionals*.

²¹² C.f. Quine’s *Methods of Logic*.

Part III: Laws and Skills – A Therapy

was a Frenchman is held fixed. In the second conditional, exactly the reverse is the case.²¹³

David Lewis takes it to be the *context* of utterance that determines which facts must remain fixed and which can go when evaluating a counterfactual.²¹⁴ Though Lewis' particular way of spelling this out – in terms of contextual determination of a closeness measure among possible worlds – is not universally adopted, the idea that the truth-value of counterfactual conditionals is essentially context dependent has found wide acceptance. Unless there is a context that sufficiently – though not necessarily exhaustively – determines what is to remain fixed and what can go, a counterfactual conditional does not have a determinate truth-value.

On the side of desiderative inferences there is something that corresponds to this context-dependence of counterfactuals. To attribute to someone's making a desiderative inference the status of being *prima facie* justified, we must also attribute to him a certain *implicit practical understanding* about *how* he is going to carry out the intention he forms with the consequent of the desiderative inference. So for someone's utterance of the desiderative inference about a match:

(DI-M) *I want to light this match; so, I shall scratch it.*

to have the status of being *prima facie* justified (that is, justified in the absence of ethical considerations and the such) it is not enough that there is a nomological connection between scratching a match and lighting it. The speaker must form the intention to scratch the match *with the implicit understanding* that he is not going to dip the match in water before scratching it, that he is going to ensure there is enough air and not too strong wind where he carries out the scratching, etc. If there is reason to believe that the speaker of (DI-M) has a habit of dipping match in water before scratching them, for example, then there is good reason to withdraw the *prima facie*

²¹³ Of relevance are also those examples of failure of strengthening the antecedent, originating from J. Howard Sobel, and reported by David Lewis, such as: If Otto had come, it would have been a lively party; but if both Otto and Anna had come it would have been a dreary party. Both counterfactuals can be true, which means that, in evaluating the first counterfactual, we must have held fixed the fact that Anna did not come to the party. That may turn out to be not possible under some circumstances, say when it is a known fact that Otto and Anna always visit parties together.

For more examples as well as Lewis' credit to J. Howard Sobel, see Lewis' *Counterfactuals*, page 10.

²¹⁴ C.f. Lewis' *Counterfactuals*, page 67, where he discusses another pair, also stemming from Quine, that is for the present purposes analogous to the pair about Verdi and Bizet.

approval that (DI-M) otherwise would enjoy. Compare the importance of implicit understanding to the *prima facie* justifiedness of (DI-M) to the role of context for the truth of a corresponding counterfactual conditional:

(CC-M) If this match had been scratched, it would have lighted.

For *each* contextually determined factor that needs to be fixed for (CC-M) to be true, there is an implicit understanding that must be presupposed for an utterance of (DI-M) to be *prima facie* justified. So if for (CC-M) to be true it needs to be fixed that the match remains dry when scratched, there must be a corresponding implicit understanding on the part of the speaker of (DI-M) that he is not going to wet the match before scratching it, in order for us to recognize his desiderative inference as justified. And ditto for all other contextually fixed factors for a truly asserted (CC-M).

In fact, an implicit practical understanding comes with *any* desiderative inferencing. Sometimes the accompanying understanding makes the desiderative inferencing unjustified, despite the existence of a suitably related nomological connection. This is exactly analogous to the fact that, despite the nomological connection between being a unlit but scratched match and being not dry, the following counterfactual is not true in the most natural context:

If match *M* had been scratched, it would not have been dry.²¹⁵

The contextual determination of which facts must remain fixed contributes as much as the availability of laws does to the truth-value of a counterfactual conditional. And just so, the accompanying implicit practical understanding on the part of the speaker as to *how* he shall carry out the intention he forms contributes just as much as the availability of laws does to the goodness of his desiderative inference.

This theme of implicit understanding and contextual determination will return again when we take a look at *skills*, the core element on the verification side of the use of law statements. This is a central feature of the linguistic practice centered around law statements, and is deeply connected, for example, with the *ceteris paribus* clauses that are felt to be attachable to most law statements. That connection will be examined in §7.3.2 below.

²¹⁵ The example is from Goodman. See his *The Problem of Counterfactual Conditionals*, page 120.

Part III: Laws and Skills – A Therapy

The practical deliberation principle for law statements is therefore methodologically equivalent with the traditional counterfactual rule.

Tying our discussions in the preceding sections all together, we should spell out the practical deliberation principle this way:

Pragmatic Side: Practical Deliberation Principle:

A law of nature statement *permissively* licenses the making of *desiderative inferences* of a suitable sort in the course of a practical deliberation.

I hope to have explained sufficiently in the course of this section what the phrases “permissively license” and “desiderative inferences” mean, and to have provided compelling grounds for thinking that this principle is a good substitute for the counterfactual rule.

§5.3 The Skill-by-Experiment Principle

§5.3.1 Skill by Experiment

The Skill-by-Experiment Principle says that one is warranted to make the law statement:

(LS) *It is a law of nature that, if any S-situation is A, then it is also B.*

if one

masters, through experimentation with S-situations, a skill to *make an S-situation B by making it A.*

This condition of warrant requires a number of clarifications in itself – that is, independent of its role *as warrant provider*. First, I assume it is not controversial that there is such a thing as *mastering a skill*. The concept of skill employed here is a pre-theoretical one. It is the same concept used when we say that a car mechanic’s work requires many skills, that Pete Sampras has an amazing skill for serves, or that most of the skills we acquire in school turn out to be useless at work. Second, I also assume that some skills are acquired by experimentation. The skill of performing simple arithmetic calculations without pen and paper, for example, might be plausibly thought of as attainable by repeated *exercise*, rather than by experiment. The same can

be said for most basic athletic skills. And arguably some other skills we have since birth, by virtue of being the animal species we are. But exercise is not the way to all kinds of non-innate skills. Athletes no doubt had to experiment a lot to develop a technique – a skill in other words – to suppress splash when they dive from a 10-meter platform. Virtually all engineering techniques are developed by experimentation. It is by experimentation that people developed the skill to communicate vocally by a connecting wire, the skill to control the direction of travel of a boat with rudder, the skill to fly, and so on so forth. Third, the kind of skills mentioned in the condition of warrant for law claims is a special kind of skills attainable by experimentation. They are skills for realizing a *complex sort of intentions* (“CI” for short) of the form:

(CI) To make *this* S-situation B by making it A.

Examples of such complex intentions include: to light this match by scratching it; to make this piece of copper carry electric current by subjecting it to electric potential differential; to increase the pressure of this chamber by pumping more gas into it. I shall now refer to this kind of intentions and the corresponding skills *binary intentions* and *binary skills*, respectively, corresponding to “binary law statements”. Fourth, although the skills in question are skills for realizing complex intentions, the *experiments* through which the skills are acquired need not begin with a schema of complex intentions. A kid might learn to light a match *not* by starting with that complex intention, but by starting with the intention to *find out what happens* when he strikes a match.

Fifth, the kind of skills talked about here require, precisely because they are acquired by experiments that require the ability to form complex intentions, presupposes at least some *linguistic* capacities. So even though the skills themselves are not linguistic skills,²¹⁶ they are not *a*-linguistic either.

The sixth point is both important and more complex, and has to do with the *implicitness of mastery* in any experimental success. The point can be appreciated by first noticing the indefinite article in “masters ... *a* skill to...”. The use of indefinite article is to accommodate the fact that, supposing the mastery of the skill in question

²¹⁶ The idea of account for nomological connections through skills can be applied also to *a priori* laws, or conceptual connections, in which case the skills in question would be *linguistic* skills.

Part III: Laws and Skills – A Therapy

is possible, there is never a *unique way* of making an S-situation B by making it A, as well as the fact that we attribute a skill even when the person has not mastered *all* different ways of doing it. But differences between different ways, or the character of any particular way of doing it need not become explicit in the consciousness of the person mastering the skill. The only *explicit awareness* required for someone to master such a skill is that, when he exercises it, he thinks of himself as trying “to make an S-situation B by making it A”. So someone who skillfully lights a match need not be consciously thinking to himself “don’t let the match get wet”, even though, keeping the match from getting wet is precisely what he *does*. What these considerations illustrate is that much of the mastery of a skill is *implicit mastery*. The function of experimentation is that it allows one to find a reliable way of carrying out the complex intention “to make this S-situation B by making it A”, though features of that reliable “way” need never be explicitly articulated in words by the one who masters it. They remain in his *motoric memory*. We shall see at the end of this chapter that implicit mastery is a correlate to *implicit practical understanding* accompanying a desiderative inference, and both are correlated with the contextual determination of what must one hold fixed when evaluating a counterfactual conditional, the correspondence of the last two of which we have just discussed. I can do no more now than flag the issue. Discussions of it in connection with *ceteris paribus* clauses of law statements and the notion of explicit making are to follow in §7.

§5.3.2 Skill-by-Experiment As Warrant

Turning now to the role of skills as *warrant* for a law claim, let me note first that the way in which the warrant obtains is not through the making of an *inference*. Neither the status of having mastered a skill, nor the experimental activities through which the mastery is achieved, is of the right category to stand to a law claim in an *inferential relation*. What this relation exactly is is the topic of §7. It suffices now to note that not all warrant for a claim is obtained by making an inference. The classical example of this is of course observation reports, which are perceptual judgments made out loud. There are on the one hand raging debates about what guarantees an epistemologically speaking good instance of perceptual experience, and how the “good” instances of perceptual experience are to be analyzed. But all parties to the debate agree that

sometimes, having (a good kind of) perceptual experience²¹⁷ is enough to give warrant for an observation report based on it.

But if skills do give warrant to law claims, as I claim, then the warrant involved is certainly not a variety of perceptual warrant. The status of mastery is not an *event*, let alone a perceptual event. Perception is indeed involved in experiments. But the activities of experimentation is more than perception, and unless *successful*, the activities as such do not provide warrant anyhow. It is important to distinguish two concepts of “success” in the context of experiments and skills. There is a sense of success that is attached to one particular, dated, experimental act. So if I experiment with lighting a match by striking it. I might be successful – in this *particularist sense* – only in my third try, and not in my first and second try. It is however another sense of success that is of direct important to us, a sense of success I shall refer to with the phrase “*experimental success*” henceforth, to be contrasted with *success of (particular) experimental trials*. Experimental success is not a feature of any particular experimental trial. Rather, it secures for the agent the status of having mastered a certain skill. To that extent, experimental success and having mastered a skill are more or less equivalent, at least for the kind of binary skills talked about here.

More will be said about the relation between experimental success and the success of particular experimental trials.²¹⁸ What I shall begin arguing in the following is that, skill-warrant is not only *not a species* of the genus perceptual-warrant, the two kinds of warrants are essentially *dissimilar*. In fact, seeing the distinctiveness of skill-warrant is critical for seeing that, contrary to what the Traditional Picture purports to tell us, the skill-by-experiment principle is *in harmony* with the practical-deliberation principle. Before coming to the logical and epistemological issues, however, I want to make two non-epistemological points that underscore just how different acquiring a skill is from having a perceptual experience or making an observation.

²¹⁷ It is important not to drop the qualification “perceptual”. For as we shall see, experimental success is *also* a kind of experience, just not a *perceptual* kind.

²¹⁸ See the section §5.3.3.3 below for more.

§5.3.2.1 The Phenomenology of Experimental Success

The first non-epistemological point is that experiments and experimental success form *experience* of a kind that is distinct from perceptual experiences. As experience, experimental success has a different *phenomenology* from that of perceptual success. The first phenomenological difference is that, a perceptual experience, when successful, is experienced as a *temporally bounded* and *singular* experience. The experimental success, however, is typically experienced over time, and the experience of it is marked by frustration, false starts, and gradualness of progress. Second, while most perceptual experiences are experienced as engaging some specific perceptual organs, the success of experiment is experienced not as mediated by any *particular* sensory organ.

Third, to see the difference in phenomenology more sharply, consider the difference of what an experimenter – say Galileo dropping balls of different weight off the top of the Pisa tower simultaneously – experiences and what the people who *look on* experiences. Both Galileo and the on-lookers *perceive* the dropping of the balls, and their falling to the ground. But it seems obvious that, while Galileo experiences a kind of *satisfaction*, the on-lookers will more likely experience, not a simple satisfaction, but an *urge to try out the same thing*. This experiential difference is also manifest when someone tries out a particular dance move on the dance floor while everyone else watches him and cheers him on, or when a child looks in amazement how his mother brings down all the domino pieces by pushing down the first one. Similar contrast exists every time we go to the circus, or watch a magic show (however with the difference that the magician and circus clowns pull off a trick they *already* mastered, rather than to experiment in order to perfect a trick). This ubiquitous contrast shows unambiguously that the phenomenology of the experience of a successful experimenter has an essential aspect that goes beyond the perceptual aspect that he shares with his on-lookers. I shall describe this additional aspect by saying that experimenter's experience has an *!it-works!* phenomenological character.

Finally, while in a perceptual experience we feel bound by *what* we perceive in which we have no say, experimental success feels *liberating*. It is an experience in which we feel our *own agency* to be active. More will be said in the concluding chapter about various aspects and the implication of this for the traditional Kantian account of experience as resulting from a receptivity and a conceptual spontaneity. The conclusion we draw now is that experimental success have an *!it-works!*

phenomenological character that is liberating, and not tied to any particular sensory organ. This experience can emerge slowly, over time, and over a long period of experimental activities, rather than be experienced in a single event.

§5.3.2.2 *Motoric Capacities & Spontaneity*

The second non-epistemological point I would like to make is that experiments and perceptions require the exercise of quite different capacities, and they result in different kinds of state. Both experiments and perceptions require admittedly the engagement of perceptual organs. Indeed, experiments partly *consist* of observations. But no experiment is possible without the active involvement of the *motoric* capacity. By contrast, the role of the motoric capacities in perceptions is a much subtler and indirect one. While it would be wrong to say that perceptions require no motoric capacities, no *actual* motoric coordination takes place in a perception. On the one hand, having motoric capacities probably constitute a *precondition* for certain kinds of perceptual experience. On the other hand, certain parts of the brain responsible for motoric activities may be active during perception. Yet, because perceptions involve no actual motoric activities, one cannot acquire *new* motoric capacities through perceptual experiences alone. Briefly, perceptual experience results in an awareness about a *particular* state of affairs, while experimental successes lead to the development of new, partly motoric *capacities*.

A somewhat different way of characterizing roughly the same distinction is to say that perceptual experiences are not *actions*, or, in Kantian terms, they are not unrestrained exercises of *spontaneity*. Experiments, on the other hand, are such exercises. In fact, as I argue in the concluding chapter of this essay, experiments and experimental success represent a striking omission in Kant's theory of experience and the recent revival of Kantianism within the tradition of empiricism, according to which the receptive character of experience prevents it from being a full act of the spontaneity.²¹⁹ Experiments are *both* experiences *and* acts of spontaneity in the unrestricted sense.

§5.3.2.3 *General Subject-Matter & Repeatability*

The last point leads naturally to the first of the logical and epistemological features about skill-warrant I want to discuss. The feature I have in mind can be put, in a

²¹⁹ For Kantianism in Empiricism, see McDowell's *Woodbridge Lectures*.

Part III: Laws and Skills – A Therapy

somewhat simplified form that can appear tendentious, this way: while experiments have *general* subject matter, perceptions have *particular* subject matter. The talk of “subject-matter” here is pre-theoretical, and is a generalization of the talk about intentional directedness. The claim is that, while a perceptual experience is directed at a particular (object or event), an experiment is directed at a *type* of constellation of objects and events.²²⁰ A sign for the correctness of this claim is that we say that the same experiment is *repeated* or replicated on different occasions, whereas two seeings cannot be literally repetitions of each other. Galileo will be performing the *same* experiment when he uses a different set of balls and moves to a different tower. But the on-lookers of Galileo’s experiments will not be repeating the *same* perceptual experience. (Their experiences might be *similar*, but that is beside the point). This suggests that the subject-matter of a perceiving is either a singular event, an individual of some other kind, or simply a particular state of affairs. By contrast, the subject-matter of Galileo’s experiment is *to have balls of different weights drop simultaneously to the ground by releasing them from the same height simultaneously*. One does not experiment with a particular, one experiments with a *type* of constellations of objects and events. Galileo’s experiments have to do with a general pattern, not with particular balls or particular letting-falls.

Putting the contrast the way I just did invites some obvious objections. First of all, we do sometimes say that an observation made by one scientist is *repeated* by his colleagues. This is a usage related to but not the same as the usage found in such statements as “The labor secretary made the same observation in last week’s meeting”. In this latter kind of usage, an observation is the *content* of a remark, rather than the objects or the state-of-affairs remarked upon. The observation repeated by a scientist is similarly not the objects or events involved. What is repeated is a *type* of constellation consisting of a control-setup of a certain *kind* together with observed occurrences of a certain *kind* under that control-setup. In fact, the talk of “repeating” an observation occurs often in contexts where the boundaries between controlled-observations and experiments are not sharp. It is therefore more likely a spill-over

²²⁰ The question about what a perceptual experience or an experiment is *directed at* must be distinguished from the question about their *content* (at least in some incipient sense of “content”). The content of a perceptual experience is a singular proposition about the object it is directed at, while the content of an experiment is a certain nomological cohesion in the constellation-type it is directed at.

effect of our talk about the repetition and replication of an experiment than anything else.

The “repeatability” criterion can be sharpened. The divergence of experiment from perceptual experience in terms of repeatability is the clearest in their respective *success conditions*. An experiment cannot be said to be successful, unless it is thought to be repeatable *with different objects* and *replicable by different agents*. It is not possible to successfully acquire a *skill* that is applicable only to one object. What would be acquired would be knowledge of “ways about” a *single* object. Similarly, no skill can be acquired that cannot be replicated by different subjects. A “skill” that cannot be replicated except by one person is not a skill, but the mysteries of a magician. Now by contrast, a perceptual experience’s success seems to be logically independent of the success of another perceptual experience of a different particular, however phenomenologically similar it is to the former. There might a sense of “must” in which we can say: if a perceptual experience, of the particular X, is successful, then it *must* be possible to repeat (by a different subject) that perceptual experience of X. But even if there is, the claim is certainly wrong if we substitute for the second occurrence of ‘X’ a Y that is not identical with X. The success of perception does not require similar success with a different particular, while the success of experiment *does* require similar success with varying particulars.

Another challenge is that it seems that people do perceive or observe *patterns*, which directly contradicts my claim that the subject-matter of a perceptual experience is particular rather than general. There are two parts to my response to this objection. The first response is methodological. The more important part of my claim is not the part about perceptual experience but the part about experiments: experiments have general subject-matter, and the same experiment can be repeated. The importance of this claim lies in the dialectical opposition I have set myself up with a particular style of philosophy: an anti-realist attitude towards laws of nature on supposedly *empiricist* ground. It is this dialectic opponent of mine, this “empiricist” philosophy, who has assumed that perceptual experience and observations can strictly speaking only be of particulars. This is in fact *explicitly* adduced as a premise in Hume’s argument against necessary connections. The “empiricist” arguments work only with another, this time silent, assumption: that perception is the only kind of experience epistemologically relevant to knowledge about laws of nature. My opposing argument consists in pointing out that *another kind of experience*, namely experiments and experimental

Part III: Laws and Skills – A Therapy

success, is also relevant (in fact, the main kind that is relevant) to epistemology about laws. Whereas I *grant* the empiricist his *explicit* assumption that perceptual experiences have particular subject-matter, I claim that at least *experimental* experience²²¹ does not have particular subject-matter. Against this dialectical backdrop it should be clear that pointing out that we sometimes speak of perceiving a (general) *pattern* in what we see is not an objection to *me*, but at most an objection to my dialectical opponent who actually relies on the particularity of perceptual subject-matter to buttress their anti-realist position about laws. My argument does not rely on that thesis. It relies on the claim that experiments have general subject-matter.

The second part of my response takes the substance of the challenge itself more seriously. That we perceive patterns or “forms” in things is taken up by followers of Kant – especially prominent are the Gestalt-psychologists – and rightly so. However, the underlying Kantian idea does not threaten a distinction between perceptual experience and experimental experience²²². Kant himself says that intuition is singular. In fact, when we say that someone sees a pattern, we mean that he sees a pattern *in such and such*, where “such and such” is a particular or particulars. That is, to say of someone that he perceives a pattern or type is to say that he perceives an individual *as* falling under that type, or exhibiting that pattern. So there are no perceptions of general aspects *per se*; there are only perceptions of general features *as being exemplified by certain particulars*.

It should be noted that I am not arguing that *singular causal connections* cannot be perceived. It has often been pointed out, I think rightly, that we *perceive* causal relations all the time: we perceive that the cat’s chasing its toy caused the milk to spill, that one billiard ball’s hitting on another caused the latter to change direction of movement, etc.²²³ It is in fact a great service of Kant and the modern Kantians to have shown us the transcendental indispensability of *general concepts* for the cognitions of *particulars*. But acknowledging that the *concept* of causality and many others are immanent in perceptual cognitions does not mean that the perceptual experiences themselves have general *subject-matter*. To put it very crudely,

²²¹ The phrase is used somewhat loosely here. What is meant is “experience of experimental *success*”. For possible misleadingness, refer to §6.3.2.3.2 below and the footnote therein.

²²² C.f. the caveat in previous footnote.

²²³ For a particularly forceful presentation of this point that is not a direct appeal to Kant, see Anscombe’s Inaugural Address Lecture at Cambridge University, 1979, *Causality and Determination*.

perceptual experiences have general *preconditions*, but they do have particular *subject-matter*.^{224 225}

§5.3.2.4 *Generality of Warrant*

The logical point about generality of subject-matter has an epistemological consequence that is crucial for the story I am telling. The point is that the warrant provided by the status of having mastered a skill is general in two senses, to be formulated in what I shall dub the “Generality of Warrant” thesis:

(GW-1) The mastery of skill provides a *direct* warrant for the corresponding general law claim, without the mediation of singular (causal) claims;

²²⁴ For more detailed discussion on the sense in which perceptual cognition can be general, especially the distinction between *transcendental* and *therapeutic* arguments, see the discussion in §6.2 below.

²²⁵ Some readers of this essay have raised questions about the concept of *kinds* that figure in the subject matter of experiments, and in the content of law-like statements. There are two questions one might ask: (i) what *sort* of generality is being claimed when I claim that the subject matter of experiments is general? In particular, is it the sort of generality based on *natural kinds*? (ii) Regardless of which concept of *kinds* underlies the sort of generality exhibited by the subject matter of experiments, how are these kinds *individuated and recognized*?

The answer to question (i) is that a successful defense against anti-realist tendencies about laws of nature will also stop anti-realism about *kinds*: the safe-guarding of the reality of laws is at the same time the safe-guarding of the reality of kinds governed by laws. However, this essay concerns an earlier stage of in the realism vs. anti-realism dialectic, and certainly does not aim at a comprehensive theory about natural kinds (or any concept of mind-independent kinds).

But with question (ii) one might be directly challenging the re-constitution of harmony through the definition of a condition of application (namely, experimental success) that requires a notion of kinds. For, if these kinds cannot be antecedently individuated, so argues the challenger, the attempt to elucidate the harmony of use for law-of-nature statements will fail. And it will fail on grounds of circularity if these kinds must be individuated by laws of nature that govern them.

This challenge fails to appreciate the dialectic force of the arguments in this essay. The idea is that challenge against the naïve realism about laws of nature as well as mind-independent kinds cannot get off the ground if the suspicion about disharmony of use is recognized as misplaced. The dialect *starting point* is naïve realism about mind-independent kinds, *inter alia*. The present essay aims to invalidate the attempt by anti-realist to shift the burden of argument to naïve realist, rather than to neutrally evaluate the realist and anti-realist positions.

Another way question (ii) can be interpreted is as a question about the finer epistemic structures surrounding natural kinds and laws of nature. An answer to it begins with the recognition that knowledge about laws and knowledge about natural kinds must grow hand-in-hand. New knowledge about laws can easily lead to revisions about what we thought we knew about natural kinds. And suspicions about a mis-classification of kinds will lead to re-formulation of laws. But that does not mean that the exploration of new laws of nature cannot be done with experiments based on the categories of kinds that we *think* we know. Fallibility in one area does not make reliance on our beliefs in that area to acquire knowledge about another area completely untenable. This is a general point about not raising the bar of knowledge too high. For a similar point, refer to section §6.4.1 below.

Part III: Laws and Skills – A Therapy

(GW-2) Not *individual* experiments, but the “general” experimental success, provides such a warrant.

In other words, the generality of warrant consists in (i) the fact that what *directly receives* warrant is a general statement, rather than propositions about particulars; (ii) the fact that what *provides* warrant is a general status, not a particular event or occurrence.

The connection between experiments and skills in the current context is that the kind of skills in question are such that they are acquired through *experimental success*. We saw that experimental success *requires* repeatability, and that experiments are directed at a *type* of object/event constellations. Because of the repeatability requirement, the mastery of a skill cannot be just a matter of a *single* experimenting act. This is why the warrant-*provider* for a law statement is not to be thought of as *particulars*, the way perceptual warrant is provided by a *particular* experience. On the other hand, because the *subject-matter* of experiments is general, what *receives* warrant from the success of an experiment is not a judgment about particulars either.

§5.3.3 Some Contrasts

§5.3.3.1 *Contrasting With Perceptual Warrant: Public Accessibility*

The contrast with perceptual warrant exists not only in the dimension of general vs. particular. So it is true enough that the warrant-provider in the case of perception is a *particular* experience, and the warrant-receiver is a judgment about a *particular*, or particulars. There is however also a contrast in terms of the *accessibility* of the warrant. It is in the nature of perception that sometimes it so happens that an object or event can only be perceived within a limited spatial-temporal region. We rely on eye-witnesses to tell us about a time long gone, a crime scene, or a far flung country. These eye-witnesses have a special authority that others do not have. That special authority is grounded in their *having seen* what they saw, and once what they saw is gone, no one else can acquire the same special authority. The warrant provided by experimental success, on the other hand, can in principle be shared. In fact, the very nature of experimental success *demand*s that it be sharable. Galileo’s claim to experimental success is *strengthened* by its repetition of the same experiment. In fact, the more numerous the repetition, the more diverse the agents that carry out the

experiment, the more varied the circumstances under which it is carried out, the more claim *we* all have to our mastery of the corresponding skill.

The point is not only that *everyone* can experience the same *!it-works!* phenomenology of an experimental success. The point is also that the replication of an experimental success not only strengthens the *original* claim by the first experimenter to his success. Replications give the *later* experimenters the same strengthened authority. Both the *phenomenology* and the *authority* of an experimental success is accessible, in principle, to everyone. So, unlike perceptual warrant, for skill-warrant there is no *special* authority enjoyed by some persons that cannot be in principle enjoyed by others. This general accessibility of authority is compatible with the fact that we do *defer* to other's authorities, both in matters of experimental success and in matters of perception. But there is a crucial difference: we defer to other's authority for experimental success for *practical* reasons, for example the fact that a person does not have the time to master *all* the skills in the natural sciences and the engineering disciplines. By contrast, deferment for perceptual authority is *out of necessity*, because sometimes we just cannot, even if we had all the time in the world, put ourselves in the spatial and temporal vicinity of the thing we want to know about.

The mandatory public accessibility of experimental success has some important consequences. First, it means that the attribution of experimental success to any *particular agent* is somewhat superficial. Ultimately, an experimental success *belongs to a whole community*, where the individuation of a "community" must be based partly on *shared skills*, and by a sense of where a replication of an experiment can be recognized as such. This sense of sharing a skill or an experimental success, based on the public accessibility of them, obviously does not militate against the intuition that, for a skill that is shared in this sense by a community, it can happen that *some* members of the community have not acquired it yet, while *some* are better at the skill than others.

The second consequence of the accessibility thesis I want to mention has to do with a kind of anti-realism different from the nomological anti-realism, which alone is a relevant target of this essay. What I have in mind is a certain *meaning anti-realism* that has come to dominate the interpretations of Wittgenstein's considerations on rule-

following. The semantic anti-realist readings²²⁶ have things backwards by basing their arguments on the implicit assumption that ultimately only individual public performances have epistemic value where linguistic understanding is concerned. The correct lesson from Wittgenstein’s passages on rule-following, instead, should be that the *status of mastering a skill* – in that context the relevant skills are of a different kind than those considered here – cannot be reduced to either *particular* performances, or facts of other sorts about *particulars*.

**§5.3.3.2 *Contrasting With Universally Quantified Claims:
Resilience & Non-Additivity***

Another cluster of features of the skill-warrant can be illustrated by contrasting it with the way universally quantified claims are justified. One way – in fact, the most fundamental way – to verify if every coin in my pocket now is a quarter is to examine successively each coin currently in my pocket. The result of each of these examinations contributes directly to the status of the universally quantified claim. If just *one* of the results turns out to be negative, the claim is disproved. For experiments and law claims, there is at first a similar dependence of experimental success, and therefore the law claim warranted by it, on a *multitude* of experimentation activities. The relation of the “many” to the “one” is however completely different. We have already noted that the many experiments relate to the law claim through the *status of experimental success*. The “generalizing step” – forgetting for a moment that this way of talking is not quite appropriate, as we shall see – occurs *before* the formation of any *claims*.²²⁷ By contrast, an examination of a coin supports directly a *singular claim*, in the form of: *this* coin currently in my pocket is a quarter. These singular claims then, in turn, support the universally quantified claim. The “generalizing step” is a step *between*, rather than before, various claims.

That it is inappropriate to think about the relation between the multitude of experimentation activities and the status of experimental success as a “generalizing step” is clearly shown by the fact that *experimental success can tolerate failures of particular experimental trials*. This is what I shall call the *resilience* of experimental

²²⁶ The most prominent ones are those of Kripke and Wright, c.f. for example, Kripke’s *Wittgenstein on Rules and Private Language*, and Wright’s *Wittgenstein on the Foundations of Mathematics*.

²²⁷ There is an *verbal* aspects to experiments, as we noted earlier. But the verbalizations involved in an experiment are not *assertions*.

success or of skills.²²⁸ The resilience shows up both *before* and *after* obtaining of experimental success. So, some attempts to increase chamber pressure by pumping more gas into the chamber might very well fail – perhaps due to issues like the airtightness of the chamber, or the heat-conducting property of the chamber walls, or some other unknown factors – yet these failures of particular trials need not preclude the eventual obtaining of the status of experimental success. This contrasts starkly with the zero-tolerance of failure in the justification of a universally quantified claim by examination of its instances. Similarly, *after* someone has achieved experimental success, occasional failure to exercise the acquired skill does not warrant rescinding the status. Someone who has the skill of lighting a match by striking it can from time to time *fail* to do just that. But we do not say, except perhaps jokingly, that, because of such occasional failure, the person has forgotten how to light a match by striking it.

The contrast with universally quantified claims goes further. For it is not just that some trial failures can be tolerated. There is also no set number of trials that one has to go through one by one – tolerating a failure here and a failure there – in order to achieve the status of experimental trial. The number of trials needed is a function of many factors, including the place, time, variety of agents who have performed the experiment, the outcome of related or similar experiments, how confident the agents feel they can reproduce the success²²⁹ of the trials, etc. It is not as if with each additional successful trial we are one step closer to obtaining experimental success. In fact, it is quite possible that after one additional successful trial, in which I notice some strange accompanying phenomenon, I am *less* certain of eventual success.²³⁰

²²⁸ I am borrowing this terminology from John Haugeland, who speaks of the resilience of skills in his “Truth and Rule-Following” in *Having Thought*, pp. 305-361. The discussion on the resilience of skills begins on page 322, section 8.

²²⁹ As will become clear in the section that follows, experimental trials cannot be judged as successful without acknowledgement that there is experimental success. So strictly speaking, to describe the additive view, we have to say “*apparently* successful trials” instead of “successful trials”.

²³⁰ This, as well as the other contrasts, does not distinguish law-of-nature statements from generic statements, nor from any other category of statements that have a law-like force. For specifically *laws-of-nature* related (linguistic and otherwise) practices, refer to the (for the subject matter) very brief discussion in section §7.4 *From Techné to Epistémé* below. However, it must be borne in mind that the purpose of this essay is to defend naïve realism about law-like connections in nature *as such*, which, if successful, is remarkable in itself. A detailed development of what distinguishes laws of nature from other law-like facts must await another occasion.

§5.3.3.3 *The Inversion of Normative Standings (& The Light that Dawns Over the Whole)*

The resilience of experimental success means that it is not achieved by *accumulating instances of successful experimental trials*. In fact, just the opposite, the “general” normative standing is more basic than the “particular” normative standings: experimental success is more basic than the success of experimental trials. To see what this means, suppose the Ideal Gas Law has not yet been established, though I suspect it is true. Suppose moreover that I try, based on this suspicion, to increase the pressure of a body of gas in a chamber by pumping more gas into it. Finally, suppose the manometer reading does go up. How do I tell whether this constitutes an instance of successful experimental trial or merely a lucky coincidence? In a way, this question is an echo of the Humean anti-realism about causality. The key to taking on this question without moving into the familiar anti-realist dialectic is to recognize that the eventual achievement of experimental success *retrospectively* enables us to tell that our first trial was a success rather an instance of luck. The key to overcoming Hume’s nomological anti-realism, as it turns out, is not a Copernican Revolution (though other considerations may still recommend one), that allegedly inverts the relationship between subject and object. The key is rather the *inversion* of the relative normative standing between the particular (trial) and the general (skill). I speak of “normative standing” because the relation between experimental success and successful experimental trials is not, strictly speaking, *epistemological*, as neither of the two are claims or claimed contents; yet the relation can be more generically described as providing support for a normative status (of success). It is the “general” status of experimental success that provides support, mostly retrospectively, for the “particular” status of having a successful experimental trial. This inverts the epistemological relation between an empirical and contingent universally quantified claim and its substitution instances: the justification of the former ultimately derives from the justification of the latter. This particular-to-general conception of epistemological authority is not merely a matter of the logical property of the universal quantifier. It in fact lies at the heart of the British empiricism to which we are heirs. Even the great Kantian amongst the empiricists, Sellars, avows his

allegiance to it with the clearest of words. As he puts it, for “empirical sentences”, the credibility or authority “flows from tokens to types”.²³¹

With “inversion of normative standing” I do not of course intend to dispute the particular-to-general epistemological order for universally quantified empirical claims. What I do dispute, is the assumption that particular observation reports are the *only* source of epistemic authority for empirical matters.²³² Another source of epistemic authority is experimental success. And here, the most basic status is a “general” one. While Sellars situates the flow of authority from the particular to the general *within* the overall participation in the game of giving and asking for reasons, without which not even a singular observation claim is possible, I am putting the same flow *along side* another source of epistemic authority: the mastery of skills.²³³

There is a phenomenon in the vicinity that is analogous to what has been described in matters of linguistic understanding with the pregnant phrase “light dawns gradually over the whole”. It is the fact that, even though one cannot establish individual experimental trials as successful without already establishing the more general status of experimental success, there is clearly a sense in which experimental success depends on individual cases that are *retrospectively* recognizable as successful trials. The “light”, here the status of success, does not come from the trials taken individually. The “light” comes to individual trials only when the issue of skill mastery and experimental success is settled. Yet *if and when* the light of experimental success shines, it shines from the experimental trials collectively, which individually would remain in the “dark”, where success and luck cannot be distinguished.

²³¹ C.f. Sellars’ *Empiricism and the Philosophy of Mind*, §32. Even though this description occurs in his characterization of a foundationalist position against which he rails, he eventually acknowledges the “logical dimension” in which this flow of authority does occur.

See the passage, now made famous by McDowell, in the section cited: “There is clearly *some* point to the picture of human knowledge as resting on a level of propositions – observation reports – which do not rest on other propositions in the same way as other propositions rest on them. On the other hand, I do wish to insist that the metaphor of “foundation” is misleading in that it keeps us from seeing that if there is a logical dimension in which other empirical propositions rest on observation reports, there is another logical dimension in which the latter rest on the former”.

²³² In the dimension in which Sellars concedes that authority flows from the particular to the general.

²³³ Ultimately, the relation between singular observation reports and the backdrop of participation in the normatively structured discursive practice is relation between particular linguistic performances and the general mastery of linguistic *skills*. So in a more complete picture, skills not only provides an epistemic access to the world *along side* perception, they also provide the background against which even the most simple observation report must be seen to be intelligible as such. For more on this, see the concluding chapter.

Part III: Laws and Skills – A Therapy

Yet another way to appreciate the Inversion of Normative Standing begins by looking at the *linguistic counterparts* to experimental trials, both successful and lucky ones, and experimental success. These counterparts are: material conditionals, singular causal claims, and law statements. Of these, material conditionals correspond to the *apparent* success of an experimental trial. Apparent success here means that I tried to make an S-situation B by making it A, and to a casual observer I appeared to be successful, even though there could be reason to think that it is merely a lucky accident that the S-situation turned B subsequent to my making it A. Singular causal claims correspond to genuinely successful experimental trials, and experimental success corresponds to a law claim. More schematically:

(i) *Apparently successful trial to make a particular S-situation B by making it A*

↔

The S-situation is A \supset the S-situation is B;

(ii) *Genuinely successful trial to make a particular S-situation B by making it A*

↔

The S-situation's being A leads to its being B;

(iii) *The experimental success to make S-situations B by making them A*

↔

It is a law of nature that any S-situation that is A is also B

Since I am only using the correspondence for purposes of intuitive illustration, I am not going to say *precisely* what the correspondence consists in. It is *roughly* the relation “carrying a piece of *information* expressed by”. So for example, a genuinely successful trial carries the information expressed by a corresponding singular causal claim. In claiming that experimental success provides warrant for law statements, and that experimental success grounds, rather than being grounded by, individual successful experimental trials, I am giving expression to the idea that law claims are *neither* established by observed material conditional truths, *nor* by singular causal claims.²³⁴ It is an old idea, one in fact championed by Hume, that singular causal

²³⁴ In particular, it is a radical mistake to think that certain kind of law statements are merely universal quantifications over singular causal claims. For a representative of such a view, see for example Wayne Davis' *Probabilistic Theories of Causation*. For a summary of arguments against such a view, see Hichcock's *The Mishap at Reichenbach Fall*, §3, and the introduction of Eells' *Probabilistic Causality*.

claims imply or presuppose general nomological truths.²³⁵ The Inversion of Normative Standing is a version of a similar idea. Later in §6.2.2.1 I shall directly argue against the view that the warrant for a causal law claim is mediated by warrant for corresponding singular causal claims. (One should however carefully distinguish the sense in which singular causal claims are prior to general nomological statements from the sense in which the opposite is true. General nomological facts are explanatorily prior, even though, given a general idea of causality, one can directly *perceive* a singular causal fact without knowing any relevant nomological facts. In that sense, then, singular causal claims can be *epistemologically* independent of specific nomological facts.)

We can now sum up our discussion about skill-warrant provided for by experimental success as follows: it is, rather than supported by individual successful experimental trials, the basis on which the success of a trial is discerned and distinguished from luck (a “light” that *dawns on the whole*); it is *resilient* in the face of occasional failures of experimental trials; it is *essentially publically accessible*, both phenomenologically and as a normative status; it provides a *general* warrant in that it directly warrants a general law statement, rather than through singular causal claims.

Before coming to the defense of the picture of linguistic use concerning law statements against various objections, as well as to its elaboration and improvement in the next chapters, we can already see how the old nagging disharmonies characteristic of the Traditional Picture have completely vanished.

²³⁵ So Hume writes: “we may define a cause to be an object, followed by another, and where all the objects similar to the first, are followed by objects similar to the second.” pp. 159-60, *An Enquiry Concerning Human Understanding*. A prominent modern proponent of the same view is Davidson. See his exposition in *Causal Relations*. For a probabilistic theoretic version of this claim, see Hitchcock’s *The Mishap at Reichenbach Fall: Singular vs. General Causation*. The relativization of singular causal claims to alternative causes (§5 of that essay), which lies at the heart of Hitchcock’s account of the relation between singular and general causal claims, corresponds to the existential generalization (over causal laws) in Davidson’s non-probabilistic version of the same Humean claim.

For counter claims, see, in addition to the article by Davis cited above, Anscombe’s Inaugural Lecture at Cambridge University, and John Carroll’s *Property-Level Causation*. Note however that these counter positions assert at best some sort of *independence* of singular causation from general causal laws. They do not claim general causal claims *depend* on singular causal claims, which is, judging from the number of proponents for it, an even more extreme position.

§5.4 The Problems of Disharmony Revisited

§5.4.1 The Problem of Induction

To see that this is indeed the case, consider first the Problem of Induction. As we saw, this is essentially the disharmony between the observation rule and the prediction rule: no finite observation of past instances ought to establish something powerful enough to support predictions about the future. Recall also the realists' argument that, in order to justify claims about future, or simply unexamined, instances based on examined instances, the justification *must* go through something that is stronger than universally quantified statement. According to this argument, solving the Problem of Induction requires, paradoxically, that one acknowledges a stronger kind of content.²³⁶ The problem with the realists, however, is that they have only got at most half of what is required to solve the Problem of Induction: they have only got the right kind of content. There remains a glaring gap between the evidence, which is about finite number of *particulars*, and a content that is *unrestrictedly general* (say a content about *properties*, for the property-relation sort of realists, for example).

In the picture I have developed in this chapter, by contrast, the observation rule has been replaced by the skill-by-experiment-rule. The proper warrant for a law statement is no longer finite. It is in fact a *general warrant*, based on the mastery of a certain kind of skill. Consequently, there is, intuitively at least, no epistemological gap to be bridged over between the particular and the general. As a matter of fact, the skill-by-experiment-rule is in *perfect harmony* with the practical-deliberation rule. The warrant for the law statement:

(LS) *It is a law of nature that, if any S-situation is A, then it is also B.*

is according to the skill-by-experiment rule:

(Warrant) *Mastery of the skill to make an S-situation B by making it A.*

The consequences permissively licensed by the same law statement, according to the practical-deliberation-rule, are desiderative inferencings of the form:

(Consequence) *“I want that this S-situation is B; so, I shall see to it that it is A”.*

²³⁶ C.f. our discussion in §4.3.2.1 above.

But clearly, mastery of the skill described in (Warrant) is sufficient to justify such steps in a practical deliberation as described in (Consequence). For, the latter is no more and no less than the verbal prelude to an *exercise* of the former.

The particular-general epistemological divide that so racked the brains of philosophers disappears from the picture altogether, at least *as an epistemological divide*. A corresponding particular-general divide exists in the picture though. On the general side of it are skills, on the particular side are *exercises* of a skill or *experimental trials*. But this is not just the same old divide in new pragmatist clothing. As our discussion above shows, the general side of the divide actually has priority over the particular side. Though it is true that experimental success and skill is acquired only after numerous experimental trials, dependence in this direction is not additive, and tolerates failures. Furthermore, the most important dependence goes in the opposite direction (i.e. the Inversion of Normative Standing): whether a trial has the *status* of success rather than of mere luck is retrospectively determined by the obtaining of the status of experimental success. The new picture is therefore not the old picture with different names attached. And the key, as far as bridging the particular-general divide is concerned, is the recognition that a *general* status – that of experimental success and skill-mastery – is in fact the more fundamental status.

§5.4.2 Presumption of Uniqueness of Meaning

Before turning to the Counterfactual Problem, let me point out a feature of the picture developed here: the warrant and consequence are *particularly tightly bound to each other*. As we have seen (Consequence) licensed is just an *exercise* of the skill whose mastery is (Warrant). Moreover, there is no room for multiple interpretation here: for the intention to exercise a skill to be justified, the subject must have mastered the skill! The consequence is not too weak for the warrant, the way that the elimination condition for ‘tink’ is too weak for its introduction rules,²³⁷ for it *requires* the warrant (i.e. the skill-mastery) to exist. This tightness of fit between warrant and consequence means that, as long as (the merely necessary condition of) harmony is satisfied, we must presume that the practice described here yield a unique meaning for each law statement.

²³⁷ See the discussion on Harmony above in §4.1.5.

§5.4.3 The Counterfactual Problem

The Counterfactual Problem was identified as the disharmony between the observation warrant and the modal- and counterfactual-consequences. The new picture I have described does allow us to maintain the old counterfactual rule without disharmony. The reason is, as we saw, a certain Methodological Equivalence between desiderative inferences and counterfactual conditionals. Specifically, of the following two:

(DI)_{t₀} *I want* this S-situation to be also B; so *I shall see to it that* it is also A;

(CF)_{t₂} If the S-situation *had* been also A, it *would have* been also B;

if the desiderative inferencing (DI) at an earlier time is *prima facie* justified, then so is the counterfactual conditional (CF) uttered at a later time (supposing that the counterfactual presupposition – that the S-situation did *not* turn out to be A – is met). So, since the practical-deliberation-rule and the skill-by-experiment rule stand in *obvious* harmony, the Methodological Equivalence of the desiderative inferences and counterfactual conditionals implies that skill-warrant is equally strong enough for counterfactual consequences. The Counterfactual Problem, therefore, is no longer a problem.

§5.4.4 Implicit Harmony

Let me conclude the exposition of the therapeutic part of my story with a remark about an aspect of both new rules introduced in this chapter. This is an aspect that will become the focus of the second part of my story in §7. This aspect has to do with, on the one hand, the fact that the mastery of skills is mostly an implicit – in the sense of non-discursive – mastery, and, on the other hand, the making of a *prima facie* justified desiderative inference requires an implicit practical understanding. I have talked about implicit practical understanding and implicit mastery in discussions of the respective principles, and promised to show that they are perfectly “matched” with each other.²³⁸

In fact, the harmony between the two principles is not completely demonstrated unless the *implicit*, non-discursive aspects of them are shown to match. Consider again the two sides:

²³⁸ This is mentioned in my discussion of the last and sixth point in the section §5.3.1 above.

(Warrant) Mastery of the skill to make an S-situation B by making it A;
 (Consequence) “*I want that this S-situation is B; so, I shall see to it that it is A*”.

Note that the implicit aspects of (Warrant) and (Consequence) need not *automatically* match up. So, someone might have, with difficulty, mastered the skill to light a match by striking it. But, he might still make the desiderative inference described in (Consequence), accompanied by the practical understanding – perhaps due to an ingrained eccentric habit of his – that he shall *dip the match in water* before scratching it. His mastery of the skill of course does not justify a desiderative inferencing accompanied by *that* kind of practical understanding. Note however that even in this kind of cases where there is a mismatch of the implicit aspects, the “disharmony” between the (Warrant) and (Consequence) is the disharmony between two different practical understandings, or two different “ways”, of realizing a certain sort of intention. This kind of “disharmony” is not of the same devastating kind as those we I have identified in the Traditional Picture, embodying as the latter do an unbridgeable epistemic divide between the particular and the unrestricted general.

In any case, though, the point of harmony does not require an automatic match of the implicit aspects. It only requires that, *given* the mastery of the skill in question, an agent has available to him *some* practical understanding of the corresponding desiderative inferencing such that (Warrant) and (Consequence) match even in their implicit aspects. But this is quite obviously true. If I master the skill to light a match by striking it, then I master a *particular way* of striking the match. My implicit mastery consists in those features of this particular “way” that is not explicitly characterized with words. But mastery of this particular “way” of striking a match – e.g. making sure the match does not get wet, ensure there is abundant air, etc – is just the kind of implicit practical understanding that I need to accompany a corresponding desiderative inferencing “I want to light this match; so, I shall strike it”, so that (Warrant) and (Consequence) would match.

The harmony, therefore, is genuine after all. But much of it is not visible to the (discursive and linguistic) eye. For much depends on a match between an implicit mastery of a skill and an implicit practical understanding of the exercises of that skill. The progress from *techné* to *epistemé*, or from simple skills to science, is precisely the making explicit of that which remains implicit in the mastery of skills. This is the topic of the next part of my story about laws of nature. Before turning to that topic, I

Part III: Laws and Skills - A Therapy

shall first consider various objections to the therapeutic proposal I have made and described in this chapter.

§6 Objections

§6.1 Introduction	243
§6.2 Perception, Empiricism, and Transcendental Arguments.	244
§6.2.1 Empiricism – “Are We Leaving The Empiricist Tradition Behnd?”	244
§6.2.2 Transcendental vs. Therapeutic Arguments – “We Can <i>Perceive</i> Causal Connections; So, We Don’t <i>Need</i> Skills”	245
§6.2.2.1 Perceptual Warrant Mediated By Singular Causal Claims	246
§6.2.2.2 Transcendental Arguments vs. Therapy	248
§6.2.2.3 Applications of Transcendental Argument Contrasted	249
§6.2.3 Perception vs. Skill – “Isn’t Perception <i>Also</i> A Kind of Skill?”	251
§6.3 Non-Perceptual Elements.....	252
§6.3.1 Laws vs. Causal Laws – “Didn’t Cartwright/Price Already Say Similar Things About <i>Effective Strategies?</i> ”	253
§6.3.1.1 A Crash Course On the History of Probabilistic Causation.....	253
§6.3.1.2 Effective Strategies	255
§6.3.2 Analytical vs. Use-Theoretic Accounts – “This Is A Manipulationist Theory of Causation In Disguise”	258
§6.3.2.1 Classical Manipulationist Theories of Causation	258
§6.3.2.2 Menzies and Price’s Agency Theory of Causation	260
§6.3.2.3 Three Differences & Two Strands	261
§6.3.2.3.1 Causal vs. Non-Causal Laws.....	261
§6.3.2.3.2 Centrality of Phenomenology	261
§6.3.2.3.3 Epistemological Access to Causation.....	262
§6.3.2.3.4 Two Reductionist Explanations: Analysis vs. Analogical Concept Formation.....	263
§6.3.2.3.5 Object Level vs. Meta-Level Conceptual Explanations.....	264
§6.3.2.4 Criticism A: Confusion of Metaphysics with Epistemology.....	264
§6.3.2.5 Criticism B: Circularity of Explanation.....	267
§6.3.2.5.1 The Charge and Menzies & Price’s Reply.....	267
§6.3.2.5.2 Color Concepts and Looking-Concepts.....	268
§6.3.2.5.3 “Acquaintance” with the Concept of “Bring About”	268
§6.3.2.5.4 Burden of Argument, Dialectical Context, and Question Begging.....	270
§6.3.2.6 Criticism C: Unmanipulable Causes	272
§6.3.2.7 Criticism D: Excessive Anthropomorphism.....	274
§6.4 Skepticisms	276
§6.4.2 Local Skepticism – “How Do We <i>Know</i> If We Have Mastered a Skill?”	276
§6.4.2 Global Skepticism – “The Account Is Circular By Assuming The Possibility of Skills”	277

§6.1 Introduction

In this chapter I shall further clarify the picture developed in the last chapter by contrasting it with positions that might appear similar to what I am doing here, and by answering objections that might be directed against what I said in the previous chapter. I formulate both kinds of reactions as objections, since for most of the issues I shall be dealing with, it is hard to strictly separate the objection from the inquiries about the relation between my account to previously existing theories.

The clarifications presented below are strategic ones and to that extent differ from those meeting immediate and local challenges that have already appeared in the course of the last chapter. They also differ from another sort of clarification that must be postponed till the next chapter. The three elements discussed in the preceding chapter – experiments, binary skills, and desiderative inferences – form only a first approximation of a stage of practice *before* the development of what we would actually term “laws of nature”. That was O.K. since our goal in that chapter was merely therapeutic: to show that the perception that there cannot be harmonious linguistic use for *any* kind of nomological statements, crude or refined, is unwarranted, and that the perception results from bad philosophical assumptions rather than having a basis in what we actually do; and to thereby remove the perpetual irritant that continues to move philosophers towards anti-realism about laws. But the picture remains an approximation, though a highly suggestive one that captures some important truths about nomological discourse. Part of a further clarification then, must be about how this crude picture can be refined and further developed to make it exhibit convincing likeness to the real practice in which laws of nature are discovered, talked about, and appealed to for various goals. This constructive sort of clarification will be developed in the next chapter.

The objections will be grouped into three sets: those about perception and the status of empiricism, those about the non-perceptual elements in the picture developed in §5, and those of skepticism.

§6.2 Perception, Empiricism, and Transcendental Arguments.

§6.2.1 Empiricism – “Are We Leaving The Empiricist Tradition Behind?”

The first objection I shall deal with is that in raising practical deliberation and experiments to the fore, I am giving short shrift to the empiricist tradition. According to this objection, the picture I delineated in the previous chapter does not sufficiently respect the empiricist credo, now the mainstream of analytical philosophy, that we depend on experiences to gain knowledge about the world around us.

Formulated this way, this objection is merely a misunderstanding. In my discussion of skill-warrant, I made it clear that experiments and experimental success are a kind of experience, and I argued in depth for the fundamental difference between experimental-experience and *perceptual* experience: phenomenologically, physiologically, logically, and normatively. The spirit of empiricism, insofar as it is something to which we should aspire, urges only deference to experience as the source and limit of our knowledge, but it does not prescribe antecedently *what kinds* of experience there are. The claim that *perceptual* experiences are the only kind of experience there are is a thesis *wholly independent* of empiricism – though assumed by traditional empiricists – so cannot be supported by the latter. Let us call the thesis that specifies the extent of experience to be exhausted by perceptual experiences the *narrow-extent thesis about experience*.

Might not the narrow-extent thesis about experience be supported by something like ontological parsimony – the principle that we should not unnecessarily multiply types of existences? There is something specious about the mentioning of such a principle in the current context. First of all, skills, experiments, and experimental success are things the existence of which is *not in dispute* at all. We are not talking about numbers, electrons, or some other contentious entities. Ontological parsimony is an appropriate consideration only when there is an ontological dispute to begin with. When there is not, to insist on ontological parsimony is to stake out an ideology, rather than to make a good argument. Second, it is hard to deny that experiments and experimental success are a kind of experience and a *different* kind from perception. My argument in the last chapter concerning their difference from perceptual experiences are all based on direct and uncontroversial observations about them and about how we speak of them, rather than based on antecedent idiosyncratic theories of mine. It is, to put it simply, a *factum*, rather than a *theoretical postulate* on

my part, that there is this different category of experience next to perception. It would be absurd to ignore the facts and, out of respect for “ontological parsimony”, to forcibly put both in one category. Experiments and experimental success cannot be assimilated to perception, and that is a fact, not a postulate.

It is true that traditional empiricists have implicitly assumed that all experiences are perceptual experiences. Because of that, and because of the immense influence of empiricism, we are not quite used to thinking outside of that assumption any more. But “habits of mind” are, as Hume taught us, not a reliable guide for truth. Another such mendacious “habit of mind” is to think that it belongs to different styles of philosophizing to make skills the central explanatory concept vs. to make experience the fundamental explanatory concept: the former is “pragmatist” philosophy and the latter is “empiricist” philosophy. This requires some explanation. The “pragmatist” philosophy that gives skills a central explanatory role is primarily associated with more recent philosophers such as Brandom and Haugeland, *not* with the classical pragmatists (i.e. Peirce and James, though Dewey might be an exception). Skills became a fashionable topic in the world of Anglo-Saxon philosophy since Gilbert Ryle’s *The Concept of Mind*, where much is made of “know-how”. In Ryle’s work, skills played a central role in the “de-substance-ing” of the mental. Brandom and Haugeland are examples of modern pragmatists of this lineage, where skill-mastery is made explanatorily basic for (linguistic or otherwise) *understanding*, which is a mental attribute or capacity. But for these pragmatists, skills never acquired a genuinely *epistemological* role. They are liked instead for their utility in *ontological demystification* of certain mental categories. The following contrasts throw into sharp relief the uniqueness of the role assigned to skills in the story I am telling: the claim of the previous chapter is, against the traditional empiricist, that *skills*, not just perceptual experiences, can have epistemological significance; and, against the modern pragmatists, that skills can also have *epistemological*, not just ontological-explanatory, significance.

§6.2.2 Transcendental vs. Therapeutic Arguments – “We Can *Perceive* Causal Connections; So, We Don’t *Need* Skills”

But, the objection might continue, granted that there *is*, next to perception, a category of experiences having epistemological significance, do we *need* it for our explanation

of laws of nature? Might it not be better to argue that we can *perceive* nomological or causal connections?

In my earlier discussion of the thesis that experiments have general subject-matters I mentioned two considerations²³⁹ that impute general content to perceptual experiences. The first is the Kantian consideration that singular perceptual experiences (“*Anschauungen*”) presuppose – are not possible without – general concepts and conceptual capacities. The second consideration was the simple observation that we directly perceive causal connections, and that to that extent, the general relational concept of causality is involved in such perceptions. Considerations like give credence to the hope that perceptions can serve as the epistemological basis for laws of nature, after all.

§6.2.2.1 Perceptual Warrant Mediated By Singular Causal Claims

Two things make the idea of basing law claims on perceptual warrants problematic, apart from the fact that no one has ever attempted to approach the epistemology of laws this way. First, although perceptions can be the source of knowledge for singular causal connections, they cannot *directly* support or warrant claims of *law*. More generally, as already mentioned in earlier discussion, the Kantian thought is that general concepts are *required* or *involved* in perceptual cognition. It is not that perceptual cognitions themselves have *general content*. Second, it is equally implausible to claim that causal law statements are supported *indirectly* by perceptions of singular causal connections, mediated by singular causal claims. The reason for the failure of indirect warrant lies in the need to appeal to some form of *inferential* warrant. For that inferential warrant for laws brings back *disharmony* in the form of the Problem of Induction again. Let me explain.

Let us examine two conceivable variants of this “indirect perceptual warrant” strategy, according to which support for general causal law statements is ultimately perceptual, but is mediated by singular causal claims. The first of these variants analyzes the *content* of a causal law statement as a universal quantification whose substitution instances are singular causal claims. According to this variant, a causal law statement can be justified by perceptually confirming each of its substitutional instances, the way “All coins currently in my pocket” can be justified by perceptually

²³⁹ C.f. §5.3.2.3 above.

confirming each of *its* substitutional instances. The second variant does not analyze the content of causal law statements in terms of singular causal statements. Quite the opposite, it assumes that what Hitchcock calls the “Humean strategy”²⁴⁰ is correct. That is, it assumes that the contents of causal *law* statements are more basic than the contents of *singular* causal statement, not the other way around. According to this variant, perception establishes a singular causal statement: *this* S-situation’s being A led to it’s being B. The singular causal statement describes an instance of *some* causal law, and so serves as evidence for the *existentially quantified* claim that there *is* a causal law claim under which the singular causal relation described falls, even though the causal law claim playing this role need not employ the general terms “S”, “A”, and “B”²⁴¹. A second perception establishes another singular causal statement, of the same form (i.e. employing the same “S”, “A”, and “B”), and a third, and so on. At some point, these singular causal statements established by perception and sharing the same form justify the hypothesis that all of them describe instances of the *same* causal law, which is formulable with the same general terms “S”, “A”, and “B”.

Neither variant is plausible. Neither of them correctly describes our *actual* practice, not even approximately. Yet the biggest problem is that both of them fail to remove the appearance of disharmony familiar since Hume. The second variant requires *projection* from observed, past instances (of singular causation) to future, not yet observed or existent instances of singular causation. That all the S-situations’ being A *up to now* has led to their being B does not logically entail that a future S-situation’s being A *will* lead to its being B. The first variant, to avoid absurdity, must construe the universal quantification it finds in every causal law statement as being over not only past, but also *future* instances of singular causation. But then it cannot be established the same way as “All the coins currently in my pocket are quarters” can be – through finite perceptual examinations – not without the kind of inductive projection that lies at the heart of the Problem of Induction.

There is a strategic point in connection with the issue of generality of perceptual experience. The discussion up to now explains why the alleged aspect of generality in perceptual experience and cognition fails to make perception an alternative to

²⁴⁰ C.f. Hitchcock’s The Mishap at Reichenbach Fall.

²⁴¹ These quotation marks are to be understood as the “Quine quotes”.

experimental success as the epistemological basis for law statements. That was the first problem: the strategy simply does not work. The second problem with this strategy is that, in its appeal to Kant, it signals a failure to appreciate the difference between the *transcendental* nature of the Kantian thought, according to which perceptual cognition has a general aspect, and the *therapeutic* goal of the present essay.

§6.2.2.2 *Transcendental Arguments vs. Therapy*

The Kantian thought to which I have been attributing the idea that general concepts and conceptual capacities are presupposed and even involved in some way in perceptual cognition is transcendental in the following sense. The Kantian thought begins by asking the necessary conditions for the possibility of intentional episodes. It does not aim to *reconstruct* or *describe* the intentional episodes in detail. It is therefore perfectly possible that the Kantian argument shows us *in the abstract* that general concepts and conceptual capacities – for example those of causality – are involved *someway or other* in perceptual cognition, yet the abstract argument does not remove concrete puzzles about how such general concepts can be legitimately employed. In the case of the concept of causal laws, there was indeed a puzzle: the appearance that our use of that concept is in disharmony. The problem is not merely that we are torn with two equally strong arguments pulling in opposing directions: the transcendental argument asks us to believe that causal nomological concepts are legitimately employed, while Hume’s discovery of the appearance of disharmony asks us to believe just the opposite. The problem is also that, because the transcendental argument is made *in the abstract*, we are more likely to be impressed by more concrete considerations and to suspect the transcendental argument defective. Our heart goes with what the eyes can most clearly see.

When there is such a situation of conflict, it does not go without saying that the transcendental side will eventually prevail. As a matter of fact, there is a recent debate in the theory of perceptual experience²⁴² that also exhibits this structure of a transcendental argument pitted against concrete puzzles. The transcendental side is represented by McDowell’s *Mind and World*, while the concrete puzzles for the

²⁴² Parties to the debate often refer to their topic as the theory of experience, simpliciter. I of course disagree, on the ground that their debate is restricted to *one species* of experience. Hence the qualification “*perceptual* experience”.

transcendental conclusion – that experiences are, in the sense intended by McDowell, conceptual – include the problem of the phenomenologically grounded fine-grainedness of perceptual experience. The conceptuality of perceptual experience debate is a case where I believe the transcendental side eventually loses out, though of course I cannot go into this debate in the context of this essay.

The case at hand is different, in that the transcendental side does, I believe, prevail. One major aim of this essay, and specifically of the last chapter, is to demonstrate that the puzzle of disharmony that undermines the persuasive force of the Kantian transcendental argument is chimerical. So the goal is not to *battle* the transcendental argument, but to support it, indirectly, by taking out the thorn – in the form of appearance of disharmony – in our philosophical eye. This is what makes the current undertaking therapeutic: it removes that which prevents us from wholeheartedly embracing the Kantian conclusion. Appealing to the Kantian thought that credits perceptual cognition an aspect of generality, on the other hand, is merely repeating the abstract transcendental argument in the face of the concrete puzzle; and repetition, like stuttering, is not known to increase the power of persuasion.

§6.2.2.3 Applications of Transcendental Argument Contrasted

It is instructive to compare the effectiveness of the Kantian thought applied to the concept of causal laws with its application to other general concepts. Anscombe, in her inaugural lecture at the Cambridge University, comments on Hume's claim that we do not observe causality with the following extremely perceptive and lucid remarks:

And when we consider what we are allowed to say we do 'find', we have the right to turn the tables on Hume, and say that neither do we perceive bodies, such as billiard balls, approaching one another. When we 'consider the matter with the utmost attention', we find only an impression of travel made by the successive positions of a round white patch in our visual fields ... etc.

– Anscombe, *Causality and Determination*, page 137.

What Anscombe here points out is that, Hume cannot intelligibly formulate his claim (that causality cannot be directly observed) without going whole hog, as it were, by denying the direct observability of *anything* more complex than colored patches in

Part III: Laws and Skills – A Therapy

our visual fields. Of course, the view proposed by Anscombe with irony here has actually been held and even been in fashion. Modern Kantians like Sellars have pointed out that the phenomenological concepts such as colored patches and visual fields cannot form an autonomous conceptual practice. They are parasitic on “objective” color concepts primarily applicable to bodies in space.²⁴³

But transcendental arguments of this sort – pointing out that seemingly minimal performances actually require as their condition of possibility comprehensive conceptual capacities – manage to restore the respectability of concepts like that of “body” and the concept of causal laws only with different degrees of success. We can make the contrast by reference to two sorts of warrants available to each kind of concepts: perceptual and inferential warrants. So on the one hand, although we have perceptual warrant for singular causal relations, we do not have perceptual warrant for instances of causal *laws*. Perceptual warrant is not available to us because, though the concept of causal laws may be “presupposed”, it does not directly figure in the singular contents of perceptual cognitions. For an *instance* of, say, the causal law concept, unlike an *instance* of the concept of bodies, is general (a causal law is general), so it cannot be perceived the way an instance of the concept of bodies can. No transcendental argument can change that fact. Since the concept of bodies and causal relations, by contrast, do directly figure in the content of our perceptual cognitions, transcendental argument may indeed help us secure the in principle availability of perceptual warrant for instances of bodies and instances of singular causal relations.

It is the lack of direct perceptual warrant that drove us to the idea of perceptual warrant *mediated* by singular causal claims in §6.2.2.1 Perceptual Warrant Mediated By Singular Causal Claims. But there we discovered that an *inferential* warrant (i.e. induction) is needed that is completely out of harmony with the kind of consequences we draw from causal law claims (i.e. the Problem of Induction). Inferential warrant for instances of laws is not available to us because of disharmony. This is, as we said, a key stumbling stone between the transcendental argument and a whole-hearted embrace of the concept of causal *laws*.

What we must conclude, therefore, is that the transcendental argument has only limited effectiveness. It is especially successful with concepts that, though

²⁴³ See Sellars’ *Empiricism and the Philosophy of Mind*, §§10, section III “The Logic of Looks”.

themselves general, have *instances* that are particulars. The concept of bodies and of causal relations are like that: an *instance* of a body and an *instance* of a causal relation is a particular. We do perceive such instances under the respective concepts. Nomological concepts, on the other hand, have *general instances*: a causal law is itself general. This fact about nomological concepts together with the *singularity of perceptual cognition* rules out the possibility of direct perceptual warrant for knowledge claims of instances of laws. Transcendental argument by itself is unable to change this. Nor is transcendental argument able to restore *inferential* warrant by removing the appearance of disharmony of use. Therefore, if perceptual and inferential warrants are the only two sorts of warrants there are – which is indeed a presupposition of the *Traditional Picture* – then no transcendental argument can help us become warranted to make law claims.

§6.2.3 Perception vs. Skill – “Isn’t Perception *Also* A Kind of Skill?”

The next objection questions how the distinction between perceptual warrant and skill-warrant can be maintained given the seemingly plausible fact that certain, if not all kinds of perceptions require skills.

The idea that some perceptions require skills need not take the form of the sometimes controversial claim that it requires the understanding of complex theories and the ability to operate complex instruments to observe, say, subatomic particles. The claim may be that even what non-controversially counts as observations requires skills in that one has to *learn* to make such observations. In particular, the ability to perceive three dimensional objects as such requires that one has learned to navigate one’s environment spatially. Similarly, babies are not born with the ability to focus their eyes or to use their two eyes in coordination to determine the distance of an object. They need about 4-5 months *practice* to be able to do these things. These facts seem to show that the most mundane kinds of perceptions are an *achievement* of practice, they are, in other words, exercises of skills learned.

The fact that perceptual cognitions are exercises of skills learned does not threaten the distinction I have drawn between perceptual warrant and skill-warrant. To begin with, the kind of skills babies have to acquire to be able to perceive distance, depth, or three-dimensionality are not the kind of skills relevant to law statements, which are *verbally structured, and binary* in the sense that they involve complex two-

step intentions. But more fundamentally, what delivers a perceptual warrant is the *successful exercise* of the skills to observe, not the *acquisition* of these skills. This is just the part of what I have termed the *generality of warrant* for experimental success and skills concerning the *providers* of warrant.²⁴⁴ Put another way, given the mastery of perceptual skills, each successful exercise of them provides a *new* perceptual warrant for an observation report. The mastery of a binary skill – the skill to make an S-situation B by making it A – on the other hand, provides *just one* skill-warrant, to a corresponding law statement. An adult will typically make many, many observation reports throughout his life. And most of them are based on the exercise of the *same* set of skills for perceptual cognition. It is not the mastery of these perceptual skills that provide warrant for each and every observation report he makes, it is his *actually seeing things* with these skills that provides such warrants.

So indeed, perceptions require skills. But the claim in the previous chapter is not that I have somehow discovered skills for philosophy. The claim is rather that a *special kind* of skills (i.e. binary skills) play a certain epistemological role in a *special way* (providing general warrant).

§6.3 Non-Perceptual Elements

The next couple of objections concern the way in which the introduction of experiments and desiderative inferences works differently from existing ideas that may superficially appear similar to one of these two elements. These objections do not question the fundamental difference between perception and some sort of practical or desiderative category. What they question is how the practical and desiderative elements in the story told *here* are any different from those that have been proposed for theories about laws of nature, and how they should work any better than the latter.

Roughly speaking the first objection points to the prominence of the concept of *effective strategies* and decision theory – both of which are closely related to practical deliberation – in debates about causation, and asks how the use of desiderative inferences here offers anything new. The second objection points to a style of explanation of the concept of causation based on *intervention* and *manipulation*, going back at least a few decades. It asks how my use of experiments

²⁴⁴ See §5.3.2.4 above.

and experimental success fails to be a continuation of the manipulationist tradition. I shall start with effective strategies and decision theory.

§6.3.1 Laws vs. Causal Laws – “Didn’t Cartwright/Price Already Say Similar Things About *Effective Strategies*?”

§6.3.1.1 *A Crash Course On the History of Probabilistic Causation*

There has been for quite a few decades now a debate about the feasibility of analyzing the concept of causation in terms of probabilistic correlations. The program began at the latest with Suppes’ *A Probabilistic Theory of Causality*. The basic idea is that causes raise the chances of their effects’ occurring. One of the key difficulties for this project is the problem of spurious probabilistic correlations (e.g. two effects of a common cause are correlated but do not cause each other). A strategy for this reductionist project was introduced by Mellor²⁴⁵, centered on an appeal to theories on *rational decisions*. The idea here is that we might, with the help of decision theory, give a characterization of the *means-end* relation based on expected utility, without resorting to the concept of causation.²⁴⁶ Once we have done that, we could then turn around to use the means-end relation to explain the concept of causation. This strategy for conceptual reduction has been referred to by Hitchcock, appropriately, as “decision-theoretic causation” theory.²⁴⁷ Whether a decision-theoretic causation theory can succeed depends crucially on the explication of means-end relation using expected utility with no appeal to the concept of causation.

As it turns out, rational decision theory based on expected utility and conditional probabilities is faced with a *similar* difficulty: that of spurious probabilistic correlations. The difficulty is classically illustrated with the so-called *Newcomb’s* problem, where an action deliberated upon is correlated with a certain outcome both *evidentially* (mediated by a common cause that is already settled), and

²⁴⁵ See Mellor’s *On Raising the Chances of Effects*. In the influential paper “Causal Laws and Effective Strategies” almost a decade earlier, Cartwright notes the structural similarity between the enterprise of probabilistic causation theory and decision theory, but she did not recommend using the latter to help with the former. In fact, she assumes that probabilistic causation theory cannot be completely successful, and argues that the reason is precisely that causal laws is needed to distinguish effective strategies from non-effective one’s in decision theory. See the next subsection below for more on Cartwright.

²⁴⁶ For a classical exposition of decision theory used in contemporary discussions, see Jeffery, *The Logic of Decision 2nd Edition*.

²⁴⁷ C.f. Hitchcock’s *Causal Decision Theory and Decision-theoretic Causation*.

causally. Depending on which probabilistic correlation is used in the calculation of expected utility, classical decision theory gives conflicting recommendations. Now, if our goal is to fix our *decision theory*, and nothing more, then it is perfectly alright to invoke the notion of *causes*. That is precisely what the *causal decision theory* does: it conditionalizes on relevant causal factors for choices and outcomes.²⁴⁸ The simpler decision theory originally formulated by Jeffery without conditionalization on causal factors is often called by contrast the *evidential decision theory*.

Obviously, a causal decision theory making explicit appeal to causes cannot be the basis of a reductionist project of explaining causation in terms of probabilistic correlations. Mellor's original suggestion was that the conditional probabilities should be calculated while holding the "causal circumstances" fixed, which is done not by appealing to causes, but by appealing to *dispositions* of objects and fields.²⁴⁹ Mellor's proposal, therefore, rests on an ontology of dispositions, rather than being strictly reductionist.

Huw Price made an alternative suggestion for carrying out the decision-theoretic causation project. In contrast to Mellor, Price's suggestion does not rest on an "extra" metaphysical assumption. The core notion introduced by Price is a probabilistic correlation's *robustness under the agent perspective*:

From the agent's point of view probabilistic relevance and causal relevance cannot diverge. To introduce an agent is in effect to assume an independent causal history to the event A [the action contemplated]. Those probabilistic correlations that survive this assumption seem to have claim to be counted as genuine effects of A.

– Price, *Agency and Probabilistic Causality*, page 169.

The idea here is not to appeal to some primitive notion of "independent causal history", but to appeal to some primitive notion of "surviving the agent perspective".

²⁴⁸ For Newcomb's problem, see Nozick's *Newcomb's Problem and Two Principles of Choice*. For discussions on it, see Gibbard & Harper 1978, Price 1986, 1991.

²⁴⁹ In Mellor's words: "What I mean is what I said in section 3: an effect's chance with a cause must be greater than it would be in the circumstances without it ... we want the causal relation [to be defined] to depend on the *actual* value of what *E*'s [i.e. the effect's] chance would be without *C* [i.e. the cause] – not on what that value would be if *C*'s absence would make a difference to it. The latter is what the phrase 'in the circumstances' is supposed to rule out. ... Basically, the causal circumstances we need are dispositions of objects and fields." (Mellor 1988, pp. 234-5, underlines added).

In Price's words, "beliefs about agent probability [i.e. what probabilistic correlations look like under an agent's perspective] are *constitutive* of causal beliefs".²⁵⁰

In an article written five years later, Hitchcock gives a more precise formulation of the assumption of "an independent causal history to" an action contemplated, which involves centrally a condition of independence of random variables between actions contemplated and relevant causal factors.²⁵¹ Hitchcock's thought was that, if we are allowed to use causal concepts to give an intuitive characterization of "agent's perspective" under which the agent's contemplated actions are assumed to be initiating a new causal history, then we can show that causal decision theory gives the *same* recommendation as the (original) evidential decision theory *based on a probability distribution defined by this fiction of agent's freedom*. That is, causal decision theory is equivalent with evidential decision theory under the agent's perspective. Since the latter does not appeal to the concept of causation, but is rather based on a primitive notion of "agent probabilities" – probabilities robust under the agent's perspective – Hitchcock's results support Price's suggestion to use an evidential decision theory with "agent probabilities" as the basis for a decision-theoretic theory of causation without circularity.

§6.3.1.2 *Effective Strategies*

Huw Price frames his suggestion for a decision theoretic theory of causation based on agent probabilities as a rebuttal of an earlier influential paper by Cartwright "Causal Laws and Effective Strategies". This is how Price situates his paper in relation to Cartwright:

To sum up: Cartwright argues (i) that causal laws cannot be reduced to laws of association, because of the problem of spurious causes; and (ii) that causal laws cannot be eliminated, because they are needed to ground the distinction between effective and ineffective strategies in Newcomb problems. In refuting (ii) we have found the means to refute (i). Agency screens off the spurious associations of a contemplated action. This means not only that there is no need for a distinctively causal decision theory, but also that we may characterize causal

²⁵⁰ Ibid., page 169.

²⁵¹ This is the condition 2) on page 521 of Hitchcock's *Causal Decision Theory and Decision-theoretic Causation*.

Part III: Laws and Skills – A Therapy

regularities as associative regularities that continue to hold from the free agent's distinctive point of view.

– Price, *Agency and Probabilistic Causality*, page 169.

So while Cartwright maintains that causal correlation cannot be reduced to statistic correlations (“laws of association”) because it is needed to distinguish effective strategies from non-effective ones, Price turns the table on her, and suggests that causal correlations can be explained in terms of something closely related to effective strategies, namely, the concept of agency, and robustness under an agent’s perspective.²⁵² Whatever the truths are in this debate, it is clear that the notion of effective strategies, or something like it, is closely bound up with the issue of distinguishing spurious correlations from genuine causal ones.

How does the role of effective strategies in this debate relate to the way the concepts of practical deliberation and experiments – especially the former – are supposed to do the therapeutic work in the preceding chapter? There are several parts to the answer to that question, but the short answer is: they have little to do with each other. The first, most important point to note is that the debate about probabilistic theories of causality is not a debate about *laws*. It is a debate about the relation between *non-causal* (“associative”) laws and *causal* laws. The notion of effective strategies is supposed to be crucial for – depending on philosophical persuasion – the intelligibility, the definition, or the explanation of this latter distinction. But it is not, at least within the confines of this debate, thought to be crucial for understanding the notion of *laws* (of nature) as such.

It would be a reasonable retort to point out that the picture I have sketched in the previous chapter is one in which *only* the possibility of *causal* laws, not the possibility of laws of nature as such, is made intelligible. For, from a skill, acquired from experiment, for making an S-situation B by making it A, only causal law claims can be extracted, but no general laws of association. For example, the ideal gas law, which describes the *relation* between three quantities – pressure, absolute temperature, and density of gas – cannot be possibly be supported by binary skills of the form discussed in the previous chapter. For the skills discussed there do not even

²⁵² That Price himself thinks of the two concepts as closely related is clear from the way he summarizes his paper in an abstract by reference to effective strategies: “I argue that provided a probabilistic theory appeals to the notions of agency and effective strategy, it can avoid the problem of spurious causes.” (underline added), page 157, *Ibid*.

have the *right form* to support the ideal gas law. They might support the claim “it is a law of nature that increased gas density leads to increased pressure of gas” (with an implicit *ceteris paribus* clause, of course). But they cannot support laws formulated in the form of *variable correlations*.

But the limitation to causal laws in the previous chapter is due to the fact that that was only the *first part* of the account. The complete account, including all kinds of laws of nature, does not envisage a different kind of evidential source for laws than acquisition of skills. The extended account will be more complicated, and will include other ideas. These new ideas include that of *explicit making*, as well as different, non-binary kinds of skills. But at the end of the day, causal laws will not be contrasted with “laws of association” in terms of the special importance of experiments or practical deliberation, analogous to the way the two are contrasted in terms of the special importance of effective strategies for the former in debates over probabilistic causation. Skills and experiments are, for the story told in the previous chapter, important for both, and do not distinguish causal laws from non-causal ones.

The second difference is this. In probabilistic causation debate, effective strategies and agency are thought to *explain* the difference, or *ground* a distinction, between causal laws and “laws of association”. The experiments and practical deliberation in the picture I am developing, on the other hand, are to provide a description that approximate certain part of our linguistic practice relevant to laws of nature, a description that makes clear that no anxiety about potential disharmony of language use, and consequently about the coherence of the very concept of law-like connections in nature, should arise. Clearly, explaining a conceptual *distinction* is a very different kind of project as showing why no doubt about the *coherence* of certain concept-use can arise.

This difference in explanatory goal is also reflected in a further difference. Because of effective strategies and agency are invoked to *explain* or *ground* a conceptual distinction, they are invoked (i) in such a way that no distinction is made between the *acquisition* and the *application* of an effective strategy, and (ii) as abstract entities. What I mean is this. Facts about effective strategies and agency that play an explanatory role in probabilistic causation theory are such as: that to X is an effective strategy for bringing about Y; that the probability of X conditioned on Y is stable under the agent’s perspective. These are not facts about one’s *coming to know* an effective strategy or the robustness of a probabilistic correlation, nor are they the

application of effective strategies or the robust probabilistic correlation. To that extent, what plays a role in probabilistic causation debate are effective strategies *as abstract entities*, not elements of human practice having to do with effective strategies. This contrasts sharply with the putatively similar elements in the picture I am developing. Here, experiments and practical deliberations are elements in *human practice*, not some abstract entities with which these practices are characterized. This of course *has* to be the case given the therapeutic nature of the enterprise: to describe a constellation of language practices surrounding law of nature statements in a way that leaves no room for suspicions of disharmony of use. The theory developed here, in other words, operates on the *meta-level* by directly aiming, not at laws or causation, but at the way we *speak about* them and *do things* related to such speech. By contrast, probabilistic causation is an enterprise on the *object level*: they are theories about how to explain or analyze the concept of causation.

§6.3.2 Analytical vs. Use-Theoretic Accounts – “This Is A Manipulationist Theory of Causation In Disguise”

§6.3.2.1 Classical Manipulationist Theories of Causation

While a meta-theory about our language practice will inevitably be about elements of that practice – in our case, experimentation and the acquisition of skills, as well as practical deliberation – a theory on the *object level* need not necessarily stay away from such elements. In fact, next to the probabilistic causation debate, there is another, somewhat related, and sometimes overlapping strand in the literature on laws of nature that does just that. For this reason, it might be more plausibly compared with the current project. This is the so-called “manipulationist” theories about causation. This is a minority tradition in the philosophy of science, and the following passage illustrates the central thought of that tradition as good as any other:

... to think of a relation between events as causal is to think of it under the aspect of (possible) action. It is therefore true, but at the same time a little misleading to say that if p is a (sufficient) cause of q , then if I could produce p I could bring about q . For *that* p is the cause of q , I have endeavored to say here, *means* that I could bring about q , if I could do (so that) p .

– von Wright, *Explanation and Understanding*, page 74

The theory was developed and defended by various authors: Collingwood (1940). Gasking (1955), von Wright (1971), Menzies and Price (1993), and Woodward (2003).²⁵³ Woodward counsels in his Stanford Encyclopedia of Philosophy that manipulationists should abandon two aspirations:

- A) *to define or describe a primitive notion of manipulation;*
- B) *to reductively analyze causation in terms of this primitive notion.*

Woodward himself follows his own counsel in his own version of the manipulationist theory. He defends the utility of a theory like his without the ambitions A) and B) on the following two grounds:²⁵⁴

- (I) Such a theory is nonetheless *non-trivial* because it yields concrete causal judgments that may be controversial;
- (II) There are *different kinds* of causal notions (such as direct cause vs. net cause) so that, even if the notion of manipulation itself is not causally innocent, employing it can still shed light on other, different kinds of causal notions.

These points are well taken. But then Woodward is simply dealing in different trade as I am: I am not concerned with setting different sorts of causal notions in relation with each other, or with coming up with a model that makes causal judgment predictions that best fit what we actually say.

More importantly, insofar as the comparison of my story to manipulationists' theories about causation is intended to present an *objection* to my story, the intention of the comparison is to charge my story with the same kind of objections that *classical* manipulationists are supposed to face. Classical manipulationists, however, are attacked precisely because of their interest in *analyzing* the notion of causes in terms of human intervention or manipulation. So what I shall do now, in order to

²⁵³ Woodward in his Stanford Encyclopedia of Philosophy counts Price's essay on Agency and Probabilistic Causation as a variant of manipulationist theory of causation. I do not think this is a good idea, for the reason that Price's essay does not really talk about *intervention* or *experiments*, but merely the notion of a probabilistic correlation that is robust under the "agent's perspective". No *concrete* interventional or experimental *events or actions* play a role in Price's essay. These become significant first in the joint paper Menzies & Price (1993) two years later.

²⁵⁴ C.f. The entry *Causation and Manipulability* in Stanford Encyclopedia of Philosophy, written by Woodward himself, last revised on Monday, October the 20th, 2008.

consider these potential objections traditionally directed at classical manipulationists, is to assume for the extent of this discussion that “manipulationists” *do* have those reductionist aspirations listed under A) and B). I shall then go through the major objections, and consider whether analogous ones apply to the story told here. A good place where these objections are collected and discussed is Menzies and Price (1993).²⁵⁵ Since it is also a major version of what I would call “classical” manipulationist account – to the extent that it does seem to have the aspirations A) and B) rejected by Woodward – it is a perfect place to contrast the sorts of responses available to me and to the manipulationists.

§6.3.2.2 *Menzies and Price’s Agency Theory of Causation*

Menzies and Price prefer to call their theory an “agency theory” of causation. The official statement of their thesis is this:

... an event *A* is a cause of a distinct event *B* just in case bringing about the occurrence of *A* would be an effective means by which a free agent could bring about the occurrence of *B*.

– Menzies & Price, *Causation as a Secondary Quality*, page 187.

The connective “just in case” is rather unpretentious, but also rather vague. It turns out that Menzies and Price had something stronger, something like *conceptual explanation* in mind:

... the central point is that the concept of causation is to be explained by relation to our experience as agents in the same way that the concept of colour as a secondary quality is to be explained by relation our experience as observers.

– Menzies & Price, page 193, *Ibid.* (Underlines mine)

What Menzies & Price mean by “experience as agents” is our experience of “doing one thing and thence achieving another” or the experience of “bringing about” something by doing another. Later in the essay, Menzies & Price speak of the experience of *success*, rather than our experience of mere *succession* (page 194, *Ibid.*). As we shall see presently, there is a confusion of two senses of the genitive

²⁵⁵ For a nice summary and discussion of these problems by an *opponent* of manipulationists, c.f. Daniel Hausman’s paper *Causation and Experimentation*.

construction “experience of...”.²⁵⁶ But at first sight at least, it seems that both Menzies & Price and I aim to give a *conceptual explanation*²⁵⁷; and moreover, in their explanation of the concept of causality, Menzies and Price invoke something similar to what I have described as *successful experiments*. But this general resemblance is misleading. The projects are fundamentally different.

§6.3.2.3 *Three Differences & Two Strands*

§6.3.2.3.1 *Causal vs. Non-Causal Laws*

To begin with, one difference I already mentioned between the probabilistic causation theorists and me continues to exist in the current context. Menzies & Price – and the manipulationists in general – direct their conceptual explanation at *causal* laws, whereas in this essay causal laws figure only as the first step towards an explanatory goal that encompasses *all* laws of nature. In fact, the central thought of the manipulationist is that what makes causation *different* from other kinds of nomological connections is the former’s deep connection to “agent’s experience”. That is, manipulationists do not propose to explain nomological connections *as such* with reference to “agent’s experience”. Just the opposite, they think that the *absence* of such a connection to “agent’s experience” is a mark for non-causal laws.

§6.3.2.3.2 *Centrality of Phenomenology*

The *reason* why the manipulationists think that *causal* laws are special for their deep connection to “agent’s experience” is revealed in their emphasis on the *experiential aspect* of intervention or manipulation.²⁵⁸ The idea is that one *experiences causality*, though not laws of nature in general. The centrality of the phenomenology of experience of success for manipulationists constitute a second difference from the

²⁵⁶ See Criticism B below.

²⁵⁷ In my case, the therapeutic project is the *first part* of an eventually explanatory undertaking.

²⁵⁸ Woodward’s “manipulationist” theory of causation departs from the classical tradition in this regard. But also partly for this reason, he no longer claims to *explain* or *analyze* the concept of causation in non-causal terms. Again, that is one of the reasons why Woodward’s variant has been set aside for this discussion.

Part III: Laws and Skills – A Therapy

story told here. I have argued that the experience of experimental success²⁵⁹ has a distinctive *lit-works!* phenomenology. But, in that story, phenomenology is just *one* of many aspects of difference, that, collectively, demonstrate the existence of a *natural kind*, if you like, of experiences, next to the natural kind of perceptual experiences. The other aspects include: difference in the capacities required for each kind of experiences, different logical characters of such experiences (their subject matter, possibilities of repetition), difference in the normative relation of such experiences to other statuses (e.g. experimental success vs. success of trials).

§6.3.2.3.3 Epistemological Access to Causation

The centrality of the phenomenology of the experience of manipulation is connected with one of two strands in the manipulationists' project. This is the strand that emphasizes the *epistemological access* through such experiences. This strand of manipulationist thinking makes the mode of epistemological access to causal laws analogous to the way we come to *perceptually experience* an object, a flower, say. It is as if a causal law – in the form of a “bringing about” relation – presents itself as an *object of experience* in what they call “agent’s experience”, the way a flower presents itself as an object experience in a perceptual experience. This is in my view not correct, and we shall come to examine this thought in Criticism B below. For now let me note the difference in epistemology between the manipulationists' and my story. what provides *warrant* for a law statement, in my story, is not an *experience* of any kind, not even the experience of experimental success, but the *status* of having mastered certain skill. That status requires not just individual’s experiences of success, but also, for example, a certain collectively shared confidence that the experiment *can be repeated* by other agents. The fact that the phenomenology of the process of skill-mastery is different from perception helps to show that mastery of skill is really a *different kind of thing* as perceiving something. But the phenomenological aspect of skill-mastery itself is *not* what provides warrant. It is the *fact* – or status – of skill-mastery that does.

²⁵⁹ Note that in this phrase “success” is crucial. The kind of experiences being discussed in these pages *cannot* be happily called “experimental experiences”. The experiences in question are experiences of *achievements*, if you like, not of any processes.

§6.3.2.3.4 Two Reductionist Explanations: Analysis vs. Analogical Concept Formation

The other strand in manipulationists' thinking concerns conceptual explanation, and is the idea that we should explain the concept of causality by relating it to the idea of bringing about one thing by doing another, of which we all have direct experience. In fact, sometimes the experiencing of the "bringing about" relation is supposed to be the basis on which the concept of causality is *formed* in the first place. One might put the idea by saying that our concept of causality is *modeled* (perhaps analogically) on the concept of agency: the idea of bringing about one thing by doing another. The formulation of von Wright is an example of this "model"-approach to conceptual explanation.²⁶⁰ A slight variant of the thought formulates the relation between the concept of causation and the concept of "bringing about" not in terms of concept *formation*, but in terms of *analysis*. So that the latter is supposed to *analyze* former, rather than playing a role in *forming* the former concept. But what makes the concept of "bringing about" fit to be an analysis, in this version of manipulationism, is also connected with the supposed fact that we *directly experience* "bringing about" relations.

To see more clearly that it is possible to devise a kind of conceptual explanation without committing oneself to conceptual *analysis* in the traditional sense of giving a definition "x is P just when ...", consider an example of this kind of conceptual explanation. Sellars has advanced the idea that some concepts can be explained as *formed by analogy with other concepts*.²⁶¹ So the concept of thought is formed by analogy to speech, the concept of electrons is formed by analogy to billiard balls, etc. Though no manipulationists have explicitly adopted this strategy, it is conceivable that they could. However that may be, the conceptual explanation envisaged takes the form of setting the concept of causation *in relation* to another

²⁶⁰ This "model"-approach is strongly suggested by, for example, the first part of the often quoted passage from von Wright, "...to think of a relation between events as causal is think of it under the aspect of (possible) action. It is therefore true, but at the same time a little misleading to say that if p is a (sufficient) cause of q, then if I could produce p I could bring about q. For *that* p is the cause of q, I have endeavored to say here, *means* that I could bring about q, if I could do (so that)p" (von Wright, *Explanation and Understanding*, page 74, underlines added).

²⁶¹ See our discussion above in §2.2.2b on reservations about Sellars' endorsement of what I have called the "pragmatist interpretation of functionalism". For more on Sellars' notion of analogical concept formation, see his *Scientific Realism or Irenic Instrumentalism*, as well as *Empiricism and the Philosophy of Mind*. In the latter essay Sellars applies the notion of analogical concept formation to explain, through his now famous "myth of Jones", the concepts of *thoughts* and *sense impressions*.

Part III: Laws and Skills – A Therapy

concept, or other concepts, where the relation is either that of analysis, or a dependence relation between corresponding concept *formations*. Either way, manipulationists have a reductionist aspiration, where the aspiration can be expressed in terms of analysis or in terms of concept formation.

§6.3.2.3.5 *Object Level vs. Meta-Level Conceptual Explanations*

This leads to the *third* difference between manipulationists and me. The kind of conceptual explanation explored in this essay is not an analysis, the way necessitarianism about laws of nature is. Nor does it attempt to illuminate the concept of causation by claiming that it is formed on analogy with other concepts, the way say theorists of analogical concept formation might do. In fact, conceptual explanation in this essay does not operate on the *object level* at all. It seeks rather to depict, on the *meta-level*, the language and other practices surrounding the concept of law. What it aims to achieve is to make plausible that there is an intelligible and *coherent practice* in which we purport to speak about laws, rather than to set the concept of laws in relation to other (object) *concepts*.

To summarize, manipulationists' project has an epistemological strand and an explanatory strand with the latter being the dominant one of the two, and the project differs from mine in at least three aspects: i) manipulationists take "agent's experience" to be deeply related to *causal* laws, rather than to all kinds of laws of nature, and they make use of this supposed fact to explain *causal* laws; ii) manipulationists makes essential use of the *phenomenological aspect* of manipulation, rather than the general status of skill-mastery; iii) manipulationists seek to explain the concept of causation on the *object level*, by relating it to other *concept*, whereas I aim to explain the concept of laws on the *meta-level*, by situating the use of that concept in a constellation of *practices* that leaves no room to suspect disharmony or incoherence.

So much for the differences. Let me now turn to the objections thought to defeat classical manipulationism.

§6.3.2.4 *Criticism A: Confusion of Metaphysics with Epistemology*

The first charge Menzies and Price take up is the claim that "agency accounts confuse the epistemology of causation with its metaphysics". They summarize the charge this way :

It is widely conceded that experimentation is an invaluable source of evidence for causal claims; the objection is that it is a confusion to suppose that the notion of agency should thereby enter into the analysis of causal claims.

– Menzies & Price, *Ibid.* page 188.

So the idea is that, say, we often come to *know* that (a particular) A caused (a particular) B by virtue of the fact that we *succeeded to bring about Bs by bringing about As*. But, according to the charge, that shows only that there is sometimes an epistemological connection between the experience of success and causal relations. It would be a confusion to infer from this epistemological fact to the bald claim that for an A to cause a B *is* for certain counterfactual about manipulation success to be true. We come to know that such and such are flowers by *looking*, but it would be absurd to claim that being flowers *is* just for certain counterfactuals about visual experiences of certain kinds to be true. Now, in their response to this, Menzies and Price do not deny that they are attempting an *analysis* of the concept of causal relation. They simply deny that the analytical or explanatory relation between “agent’s experience” and causation is not the consequence of a principled verificationism about concepts, but reflects the *specialness* of the concept of causation. Menzies & Price’s point is that one need not be a verificationist to recognize the dependence of the color concepts on certain kinds of visual experiences we have. Recognizing such “response” dependence is to recognize a special feature of *color* concepts, rather than to adopt an across-the-board crass verificationism about *all* concepts. Just so, one need not be a verificationist to recognize a “response dependence” of the concept of causation on “agent’s experience”. This is how they put it:

To explain our point, let us turn to the colour analogy. For definiteness we have settled on the dispositional theory of colour, according to which an object is red, say, just in case it would look red to a normal observer under standard conditions. This theory makes colour a secondary quality in the sense that the concept of colour is taken to be an extrinsic or relational one, where the constitutive relation is to a certain kind of human response: in the case of colour red, the ‘looks red’ response...We propose that the agency approach to causation should be seen in the same light...the concept of causation is to be explained in terms of the way in which an agent’s producing, manipulating, or ‘wiggling’ one event affects the probability of another event. It is apparent that

Part III: Laws and Skills – A Therapy

this analysis has a bearing on epistemological issues... But once more it would a serious mistake to suppose that this account depends on a verificationist fallacy.

– Menzies & Price, *Ibid*, pp. 192-3.

Setting aside the question whether it is appropriate to appeal to an analogy of the concept of causation with color concepts, it is quite clear that the story told *here* cannot be the object of the same charge. The reason is that the *point* of the story is not to provide an *analysis* of the concept of causal laws at all.

Still, given the way I have motivated my story, there is a risk of guilt by association. For I have borrowed the idea of harmony, the key concept for the therapeutic project here undertaken, from Dummett, who not only is *the* most prominent verificationist in modern times, but also developed the idea of harmony in a lecture series devoted to his antirealist and verificationist views.²⁶² To preempt any suspicion of guilt by association, we need to recall both the differences between the version of harmony requirement I have adopted and Dummett's version, as well as Dummett's additional commitments unrelated to harmony. So first, instead of formulating a *sufficient* condition for the inferential coherence of the use of a concept, I have opted for a merely *necessary* condition. Second, much of Dummett's lecture series was devoted towards a *reductionist* project: to "justify" the total inferential use of a logical constant on the basis of the *verification side of use* alone. Third, even this "reductionist" tendency in Dummett's discussion of logical rules, which I do not share, is only *indirectly* related to his verificationism. The supposed impossibility to "justify" the classical rules governing the use of negation-sign on the basis of its introduction rule alone is the *ground* Dummett provides for his rejection of the classical negation.²⁶³ And it is the rejection of classical negation that is supposed to directly support a verificationist view about meaning. As already mentioned in the chapter on harmony, I share neither of these two commitments: I do not believe it is helpful to think of the pragmatic side of use rules to be "already contained" in the

²⁶² I am referring to the lecture series out of which *The Logical Basis of Metaphysics* grew.

²⁶³ See, for example, the passages on page 299 of *The Logical Basis of Metaphysics*: "This more detailed look at classical negation confirms what we had already concluded, that it is not amenable to any proof-theoretic justification procedure based on laws that may reasonably be regarded as self-justifying. ... Intuitionistic logic, however, has come out of our enquiry very well." The notion of "self-justifying" is that a set of rules can "justify themselves" as completely determining the meaning of a logical constant, if the pragmatic side is, roughly speaking, the maxim of what can be justified – in the reductionist sense mentioned in the main text – on the ground of the verification side, and vice versa.

verification side of use, nor do I endorse an intuitionistic logic. More to the point: these commitments are *irrelevant* for the therapeutic project here.

§6.3.2.5 Criticism B: Circularity of Explanation

§6.3.2.5.1 The Charge and Menzies & Price's Reply

The charge of circularity is that the concept of “bringing about” or manipulation is itself a causal concept, and consequently, by relating the concept of causation to it fails to explain causation. Menzies & Price’s response is that the concept of “bringing about” does not need to explain by appealing to causation in turn. For “bringing about” is a concept with which we have “non-linguistic acquaintance” in the experience of success, and consequently needs no *verbal* explanation, let alone verbal explanation in terms of causation. This is how they put it:

The key to seeing that this theory [of colour] is not circular is to recall that colour terms, like the terms for other secondary qualities, can be introduced by ostension. Thus a novice can be introduced to the concept ‘looks red’ by being shown samples of red: the salience of the redness in the samples and the novice’s innate quality space should suffice for him to grasp the fact that the samples look alike in a certain respect. ... The dispositionalist [of colour] can explain the concept ‘looks red’ by ostensive definition, without having to rely on any colour concept. A similar story may be told in the agency case. ... We might say that the notion of causation thus arises not, as Hume has it, from our experience of mere *succession*; but rather from our experience of success: success in the ordinary business of achieving our ends by acting in one way rather than another. It is this common and commonplace experience that licenses what amounts to an ostensive definition of the notion of ‘bringing about’. In other words, these cases provide direct non-linguistic acquaintance with the concept of bringing about an event; acquaintance which does not depend on prior acquisition of any causal notion. An agency theory thus escapes the threat of circularity.

– Menzies & Price, *Ibid.*, page 195 (underlines mine).

§6.3.2.5.2 *Color Concepts and Looking-Concepts*

To see how good the reply is, consider first the case of color concepts. Sellars, in his *Science and Metaphysics*, points to the danger of a systematic confusion in talking about experience: between experiences requiring no conceptual capacities and those that are already conceptual episodes.²⁶⁴ For experiences in which the concept of “looking red” is allegedly “ostensively” definable, the corresponding question is: do such experiences already involve an *application of the concept* of “looking red” or not? If, on the one hand, the answer is “no”, then the concept of “looking red” is only applied by us rich concept users to *describe* such experiences. But this fact, that rich concept users can describe certain experiential episode with the concept of “looking red” has no tendency to show that looking red is *conceptually prior* to *being* red, so that it can serve to *analyze* the latter. If, on the other hand, experiences in which “ostensive definition” occurs *does* involve the application of the concept of “looking red”, then the subject must already *possess* the concept of “looking red”. Sellars argues persuasively that it is not possible to possess the concept of “looking red” without already possessing the concept of being red. The idea of the “ostensive definition” may however involve, not directly the concept of “looking red”, but the *deictic* concept “looking *that* color”. Some authors, notably McDowell, have argued that there can be experiential episodes in which the deictic concept “*that* shade (of color)” is both formed and “applied” in some sense.²⁶⁵ But McDowell does not argue the point for *appearance* concepts like “looking *that* color”. Doing so would be, in Sellarsian parlance, committing to the “myth of the given”.

§6.3.2.5.3 *“Acquaintance” with the Concept of “Bring About”*

We need not get to a definite conclusion about color concepts. The little depth I have gone into the issue is helpful for bringing out the structure of the case of causal laws à la Menzies & Price. The idea that we are “acquainted” with the concept of “bringing about” in experience of success, through some sort of non-linguistic “ostensive definition”, is supposed to be understood in analogy with the case of color concepts.

²⁶⁴ This is the topic of the first chapter of that book §1 *Sensibility and Understanding*. Sellars finds this confusion in Kant’s notion of *Anschauung*. The point in the main text does not depend on this controversial exegetical claim.

²⁶⁵ See his *Mind and World*, pp. 56-60.

So we can ask the same question here: do those episodes of the experience of success involve the *application* of the concept “bringing about” themselves, or are they merely to be *described* with that concept? Like the case of colors, there are two possibilities, only one of which is possibly compatible with Menzies & Price’s story. The scenario not compatible with their story is when these experiential episodes do *not* involve the application of the concept of “bringing about”. For the fact that *rich* concept users must describe these experiential episodes with the concept of “bringing about” does not imply anything about the priority of that concept vis-à-vis the concept of causation. The scenario compatible with Menzies & Price’s story is when in these experiential episodes the subject *applies* (perhaps simultaneously *forms*) the concept of “bringing about”.

Unfortunately, the scenario compatible with Menzies & Price is not plausible at all. As a concession to them, note first that there are indeed cases where the subject, in her manipulation-experience, formulates to herself a goal *verbally* – “I want a B” – and then searches for a means by asking herself “what would *bring about* a B?” or by affirming to herself “an A will be/is needed to *bring about* a B”. So it seems that at least in *some* cases, the subject’s experience involves applications of the concept “bringing about”. But that is not the sense of “application” needed. First of all, we need the subject to get her first acquaintance of the concept through experience. Deliberating about means by asking “what would *bring about*...” requires the subject to have *already* mastered the concept, and does not represent an initial “acquaintance” with it. Second, the verbal aspects of manipulation are not applications of concepts in the *representational* sense. The experience of successfully bringing about B by doing something else may *presuppose* the ability to verbally formulate goals, but that experience as such does not purport to *represent* anything in the world.²⁶⁶ It is at this point that close scrutiny of the genitive expression “experience of ...” becomes crucial. In *one* sense, a successful manipulation is indeed an experience *of* bringing about one thing by doing another. This is the sense that, the successful manipulation is an event that can be *described* in terms of the agent’s bringing about one thing by doing another. But in *another* sense, the genitive expression refers to an experience

²⁶⁶ This is also why the story told here is not strictly *inferential*: experimental success is not a representational occurrence or act, so cannot stand in inferential relation, strictly so-called, to another claim. What is needed is the concept of “explicit making”, which I develop further in the chapter that follows.

whose *representational object* is the bringing about one thing by doing another. In that sense, it is simply not true that a successful manipulation is the experience of bringing about one thing by doing another, for the simple reason that a successful manipulation is not representational at all. The genitive expression has therefore one *extrinsic, opaque* sense, and an *intrinsic, object-of-representation* sense.²⁶⁷ Menzies & Price’s strategy of “ostensive definition” can only work if they equivocate on these two senses: it can only work if successful manipulation is an experience *of* bringing-about in the object-of-representation sense, though the sense in which it is plausible to say that is the extrinsic sense.

Note that denying that the experience of success involves any representational application of the concept of “bringing about” is compatible with the thesis that such experiences are *important* for forming the concept of “bringing about”. The point here is that we ought not exaggerate the role of such experiences. We do not form of the concept of “action”, for example, simply by acting. We need to develop a *third-person* perspective on actions to count as having the concept. If we were “acquainted” with the concept of action simply by having experienced what acting is like, all small babies would be “acquainted” with the concept of action. But that seems absurd. Overall, it seems that manipulationists overestimate the role of the *phenomenological* aspect of manipulative success for the *formation* of the concept of causation or causal laws.

It turns out that, while there is *some* parallel with the case of color concepts, the story about causal concepts is *more* problematic because a) successful manipulation, *unlike* color perception, is *clearly* not representational, and b) the formation of action-related concepts, including “bringing about”, requires, again *unlike* the formation of color concepts, a third-person perspective.²⁶⁸

§6.3.2.5.4 Burden of Argument, Dialectical Context, and Question Begging

By virtue of having a project that does not seek to make object-level conceptual explanations, the story told here is not forced to find a “rock-bottom” concept on

²⁶⁷ This distinction is inspired by Sellars, who discusses the various senses of “sense impression of...” in, among other things, the first chapter of *Science and Metaphysics*.

²⁶⁸ It would make no sense to distinguish the first person perspective from the third person perspective for *color* concepts.

which to build an explanation. I freely acknowledge that the concepts of experimental success, skills, and practical deliberation, are *all* “implicitly” nomological. In fact, since I maintain that experimental success is a sufficient *condition of application* for the concept of law, I must be committed to the phenomenon of experimental success’ being somehow nomological. But I am not forced to establish the possibility for acquiring these concepts *independent of* and *prior to* acquiring the concept of laws of nature. For I do not attempt to explain the latter by relating it to these other concepts, either by analysis, or by analogy. That would be to attempt the impossible task of, to paraphrase Brandom’s turn of phrase,²⁶⁹ trying to cook laws out of non-nomological stews. I relate, not the *concept*, but the *practice of using* the concept of laws of nature, to the *practice of doing* experiments, *acquiring* skills, and *deliberating*, and show that the way these components relate to each other is harmonious. One can *in practice* have experimental success without possessing the *concept* of experimental success.

It is true that my meta-level account assumes the possibility of genuine nomological phenomena, including skills and experimental success. But that does not beg the question if we remember the *dialectical context* in which the story is told. The issue of “question-begging” largely depends on who has the *burden of argument*. The story here is not meant to get to nomological concepts by “bootstrapping”, assuming a starting point where the *realist* has the burden of argument. The fact of the matter is, naïve realism about laws of nature – that they exist – is the pre-theoretic position. So any challenge to naïve realism must take up the burden of argument. As we have seen, one way of doing that goes back to Hume, and it begins by pointing out a supposed disharmony of language use connected with law statements. Since the anti-realists have successfully made it appear that the disharmonies really *do* exist, they appear to have successfully shifted the burden of argument back to the realists. The story here must be seen against this background. It is aimed at the starting point of this skeptical argument based on perceived disharmonies of language use. The point of the story is to nip the skeptical dialectic in the bud by showing that the *first step* of that dialectic contains an error: there *is* no disharmony of language use. The goal, in other words, is to show that the anti-realist is *not successful* in his attempt to shift the burden of argument: he still has it. At the point of this first step – pointing out the supposed disharmony – the anti-realist has not yet reached his anti-realist conclusions, and

²⁶⁹ Brandom’s original point is about normativity, not about nomological concepts.

continues to have the burden of the argument vis-à-vis naïve realism. To successfully argue that this first step embodies an erroneous assumption, therefore, does not require that we avoid appealing to naïve realism about laws. For until the anti-realist successfully develops his skeptical dialectic, it is *he* who has the burden of argument. To refer to nomological concepts in this context, therefore, cannot be “begging the question”: the anti-realist needs his first step to be correct to *earn* the right to that skeptical “question”. While the story here is precisely aimed at showing that he is *not* successful in his very first step, and hence *not* successful in shifting the burden of argument back to the realists.

§6.3.2.6 Criticism C: Unmanipulable Causes

The main reason that the first two common criticisms against the manipulability theories of causation do not apply to my story is because the explanation I venture to give is not one of *analysis*. What makes my explanatory strategy less vulnerable to the next charge is somewhat more complicated.

The criticism is that there are causal relations that cannot be manipulated, effects cannot be brought about by humans, such as the effect that fusion reactions on the sun produce light and heat on earth, or earthquakes caused by the movements of tectonic plates.²⁷⁰ Roughly, the thought is that, in order for the counterfactual in Menzies and Price’s analysis of causal relation, we need to require at least that the antecedent be *possible*. But in cases such as the ones just mentioned, the causal antecedent is something we *could* not bring about. Menzies and Price attempt to meet this charge by saying that situations such as these share certain “basic intrinsic properties” with situations we *can* manipulate. So the analysis of causal relation gets modified to be “a pair of events are causally related just in case the situation involving them possesses intrinsic features that *either* support a means-end relation between the events as is, *or* are identical with (or closely similar to) those of another situation involving an analogous pair of means-end related events” (page 197, *Ibid.*). Critics point out, rightly, that for this analysis to work, the intrinsic features shared have to be explained in causal terms.²⁷¹ But then the claim is no longer good as an *analysis* of causal relations.

²⁷⁰ C.f. Menzies and Price, page 195, and D. Hausman, page 145.

²⁷¹ C.f. e.g. Woodward, *Causation and Manipulability*.

Contrast the situation with my use-theoretic account of laws. Since I do not venture to give a theory about singular causal relation or explanations based on singular causal relations, unmanipulable causal relations do not directly affect my account. But there may be a similar concern about areas where we cannot freely experiment in order to acquire certain skills, but still would like to make law claims. Typical such areas include evolutionary biology, planetary movements, and, more generally, cosmology. In those areas it is arguably the case that observation is all we in principle can have. I believe that the status of law claims in areas such as these two sciences is in fact somewhat derivative of law claims in areas where we *can* directly experiment. But there are other ways for us to justify, indirectly, these law claims. On the one hand, the basic laws of physics, chemistry, etc. continues to hold in evolutionary biology or cosmology. The reason is that, the *warrant* for these fundamental laws does not consist in *singular* cases of possible experimentation, but the *general* status of having mastered a certain skill.²⁷² On the other hand, given that our general experience with nature is such that we keep discovering lawfulness where we can play, manipulate, and experiment, we are justified to assume there *are* laws in areas we cannot properly acquire skill-warrant. It is against the backdrop of this assumption that we infer laws from *merely observed* regularities: regularities that we cannot possibly reproduce by ourselves. Here we might say that *induction* has an application after all. It has an application first in the sense that we infer, inductively as it were, from the *existence* of laws for which we have proper skill-warrant to the *existence* of laws for which we cannot have such warrant, and secondly in the sense that we use induction to arrive at law claims for which we have no proper warrant based on *merely observed* regularities. But the application of induction in these ways do not generate the kind of vicious circle Hume is famous for pointing out. The reason is that we do not come to know *every* law by induction.

Another way the limitation of the human capacity to bring about effects in nature can generate a worry about my story has already been implicitly mentioned, and it concerns the way we conceive of our law claims to apply to *any* concrete situations, even to those where we cannot actually *exercise our skill* on the basis of which the law claims are warranted. To properly appreciate this worry, we need to distinguish it from the related worry that, for example, many laws – for example the

²⁷² This continues to be the case after the modification/extension to be undertaken in the chapter that follows.

law that increasing the density of a body of gas also increases its pressure – can *break down* in extreme circumstances (in the case of the law about gas, it breaks down when the density is extremely high so that the “gas” is no longer gas, so does not have pressure in the sense applicable to gas). The proper way to reply to the worry about extreme breakdown conditions is to say, as I shall expand in the chapter that follows, that a) extreme conditions remain *implicit* in the kind of binary skills discussed in the foregoing chapter, and b) more precise law statements including explicit clauses about the exceptional cases – including the extreme breakdown conditions – are obtained by *making explicit* what was implicit in simple binary skills, by expanding these skills into a more complex ones.

The worry from limited human capacity is however a different worry. It is the worry that, the law about gas mentioned above, for example, might not apply to bodies of gas that i) though in conditions that are “normal”, yet ii) are temporally and spatially so far from our entire humanity that it is not possible for any of us to exercise the skill of increasing their pressure by increasing their density. *This* worry is however misplaced. For the worry presupposes that it makes sense to ask whether we *really* mastered the skill to increase gas pressure by increasing gas density, or whether we merely mastered the skill to do that for bodies of gas *within some spatial and temporal confines*. Apart from an unreasonable skepticism about skill-mastery, to which I shall turn shortly below,²⁷³ that distinction exhibits a conceptual confusion. For, while there can be questions about whether we *really* mastered a skill, it does not make sense to speak of a skill for doing such and such only in situations *within certain temporal and spatial confines*. The notion of *skills* applied here is a *general* one.

§6.3.2.7 Criticism D: Excessive Anthropomorphism

The last of the four major criticisms against manipulability theories of causation is that causal relations so understood will be too anthropomorphic and cannot support our intuition that there are *objective* facts as to what causes what. Now this charge appears especially severe for Menzies and Price’s variety of agency theory of causation. The reason is that, as they themselves emphasize, the notion of causal relation is, according to their account, an *extrinsic* one, to be explained by reference to

²⁷³ See the section on Local Skepticism.

something (i.e. a potential experience of success by a subject) other than the objects involved in that relation. This means that the concept of causal relations is perilously close to the concept of tastes, for example. Which notions of tastes there are depends on certain accidental physiological features of ours. Menzies and Price argue that agency is a feature that a creature either has or does not have at all, so their account of the notion of causation does not yield a concept that is susceptible to any significant variation depending on features of the creature referenced. There is a difference between tastes and causation because, although there are many possible constellations of tasting buds we might have had, there are not different *kinds* of agency we might have had.

While this observation about the difference between agency and tasting buds is both true and insightful, my account of law statements is not even vulnerable to a similar charge of anthropomorphism to need such finessing maneuvers. The reason is that, again, I am not proposing an *analysis* of law-claims in terms of experiments. I do not believe that a law is to be analyzed by reference to something else. Experimental success is the *way* we come to know about law-like relations, and as such, may very well be unique to humans. In what other sense would laws be excessively “anthropomorphic” due to our experimental access to them? We would be committing to excessive “anthropomorphism” if we were to say that a law-fact *is* the fact that *we* have acquired through experiment some reliable skill. But that is plainly as absurd as saying that a flower’s blossoming known through perception *is* the fact *we* or whatever agent or agents have perceived it to blossom. The same charge would persist even if we sprinkle the analysis with counterfactuals and modals. Treating elements of an inferential practice as *analysis* of each other is a confusion that leads to *idealism* in the case of perception, and to (a genuinely harmful) *anthropomorphism* in the case of experiments.²⁷⁴ But unlike the manipulationists, I have not made the claim of analysis. Second, to some degrees, our biological makeup, the specific human way of interacting with nature is going to shape the way claims of knowledge are *formulated*. So anthropomorphism may get into laws by way of certain concepts we possess because of our unique biological makeup. But note that this sort of

²⁷⁴ Philosophers who advocate some version of manipulability account of causality often realize that, and make a distinction between an analysis and an account of the evidential structure of causal claims. For example, see Huw Price, *Agency and Probabilistic Causality*, part 7. The discussion is brief and half-hearted. And it is not clear what the *point* of the envisaged epistemological account of causal laws in terms of agency would be.

anthropomorphism, arising from the level of *concepts*, is shared by *both* law-claims and by claims justifiable by perception alone, and it does not infect the objectivity of *truth*. Whether these claims, once formulated employing perhaps distinctly human concepts, are in fact true or not, is something in which nature has a say. Being formulated does not guarantee being true. We are willing to grant that status only if the claims have passed the experimental test or are backed up by perception.

§6.4 Skepticisms

§6.4.1 Local Skepticism – “How Do We *Know* If We Have Mastered a Skill?”

Can we ever *know* that we have mastered the skill to make an S-situation B by making it A? Is there not always the possibility that, after some apparently successful experimental trials, we discover that these are all happy coincidences? This sort of skepticism may initially sound pressing, but a little thought reveals that it is just a corresponding version of the bald Cartesian skepticism concerning *perception*. The perceptual version of the skepticism asks how we are justified to believe in *anything* we perceive, given the fact that it is *always* conceivable for an apparently genuine perceptual episode to turn out to be a case of hallucination or whatever. It has now been generally recognized by philosophers that this sort of skepticism assumes too high a bar for knowledge claims based – in the case of the traditional Cartesian version – on perception. The similar skepticism against knowledge about skills is, similarly implausible.

The point can be put by saying that, claims about experimental success, like claims based on observations, are generally *defeasible*. Observation as evidence for a claim can be defeated by justified concern that the observation was done in non-standard conditions. Similarly, claims to experimental success can also be defeated by evidence that the agent was not completely in control of what he was doing – he may fall short of normality as a practical agent – or that there were other, unaccounted for factors that might have influenced the outcome of the experiments. For observation, it is possible to check against defeasors by “having another look”, examining more closely the circumstance of observation, and by comparing with other sources of information. For experimental success, defeasors can be checked by repeating the experiment in question with different agent and controlling for the suspicious, possibly unaccounted-for factors. In any event, there is no reason to conclude, from

the general *possibility* for observations to be defeated, that *no* observations are sound. Just so, the general possibilities for experimental success to be really “flukes” ought not generate an overall skepticism about *all* claims to experimental success. Defeasibility is not an extra price one has to pay to buy into the story told here. It is a perfectly general character of human knowledge.

§6.4.2 Global Skepticism – “The Account Is Circular By Assuming The Possibility of Skills”

Setting a high bar for knowledge claims is *one* way to generate skepticism. The resulting skepticism was a “local skepticism” because it was directly aimed at one particular part of the story told here: the part about having experimental success. Another way to generate skepticism – a *global* one – is to ask what is achieved by the *entire story*. According to this global skepticism, to tell the story, I must presuppose the very possibility of skills and experimental success. But to presuppose that is to *beg the question* vis-à-vis the anti-realists about laws of nature. For anti-realists about laws of nature will also deny the possibility of *genuine* experimental success, precisely on the ground that there are no real nomological relations in the world, and so, *a fortiori*, no such relation between an agent’s action and its effect. The skepticism, in other words, is a variant of the circularity charge already discussed.

I have already dealt with the circularity charge by pointing out the dialectical function of the account offered here (c.f. §6.3.2.5.4 above). The point was that I could not possibly be “begging the question” vis-à-vis the anti-realist when I am trying to show that the first of the steps needed for the anti-realists to *earn* his right to the “question” is in fact an error, that is, when I am showing that he does not succeed to shift the burden of argument back to the realists.

Part IV: Beyond Therapy

§7 Making Explicit & Epistemé

§7.1 Introduction	280
§7.2 Law Statements Make Explicit Skills	282
§7.2.1 The Dialectical Division of Labor	282
§7.2.2 Interlude – The Context Sensitivity of Semantic Explanations	283
§7.2.2.1 Van Fraassen & The Pragmatics Of Explanation.....	283
§7.2.2.2 Semantic Explanations.....	284
§7.2.2.3 Identity Explanation & Functional Roles	286
§7.2.3 An Identity Explanation	288
§7.2.4 Explicit Making As A Fabric Of Functional Explanation.....	289
§7.2.4.1 The Notion Of Proto-Content	291
§7.2.4.2 Nomological Connection Between Means and End	294
§7.3 Implicit Mastery of Techné	295
§7.3.1 Implicit Mastery.....	296
§7.3.1.1 Implicit Responsiveness to Degrees.....	296
§7.3.1.2 Implicit Differential Responsiveness to Different Factors	297
§7.3.2 <i>Ceteris Paribus</i> Clause of Binary Laws	298
§7.4 From Techné to Epistemé	300
§7.4.1 Justifying vs. Explaining Scientific Laws	300
§7.4.2 Warrant	301
§7.4.3 Systematicity of Science	302
§7.4.4 Science as Expression	304
§7.4.5 Absoluteness of Scientific Laws.....	305
§7.4.6 Empirical Conceptual Innovations	305

§7.1 Introduction

With a formulation from Aquinas, and more recently, of Sellars, we can say that skills are first in the order of knowing, but laws (as discovered²⁷⁵ by science) are first in the order of being.²⁷⁶ The sense of the latter claim is that, if the world had not been lawful, it would have been impossible to acquire any skills. Skills presuppose, metaphysically speaking, the presence of laws in nature. The sense of the former claim is *not* merely that there are communities, or stages of civilizations for which the concept of laws of nature was not known, whereas civilization itself is not conceivable if skills were not learned, taught, and attributed. The sense in which skills are first in the order of knowing is above all that, we come to *make law claims* by first getting into the practice of acquiring and attributing skills. Without the *practice* of enjoying and attributing experimental success and practical deliberation, we would never come to *justifiably make knowledge claims* about laws.

In §5-6 I have described how this law-claim enabling practice looks like, but only for a particularly simple form of laws: the so-called *binary* laws. In this chapter, I develop the picture of §5 further by expanding it into a fuller and more realistic account of our law-related practices. The transition from §5 to the present chapter is best thought of as a characterization of the sort of additional practices that have to develop for a community to not only count as mastering a variety of *technés*, but also to count as mastering *epistemé*, in something like the Aristotelian sense. In more concrete terms, the transition is to enable us to see how we, *in the order of knowing*, come to justifiably claim, *not only* simple, two-termed, non-numerically formulated, binary-skill based law statements with attachable *ceteris paribus* clauses, *but also* arbitrarily complex, often mathematically expressed law statements mostly thought as absolute, which are based on far more complex skills. The additional practice, in other words, must enable us to move from knowledge of binary laws of the form:

(LS) *It is a law of nature that, if any S-situation is A, then it is also B.*

to knowledge of laws such as:

²⁷⁵ I am using this in my view tendentious term without intending a judgment on whether laws are discovered or made by man. In fact, the truth is complicated. The laws arise in the acquisition, and refinement of skills.

²⁷⁶ Recall our discussion of Sellars' view on the relation of thought and language, in §2.2.2c.

$$(1) \quad F = Ma; \quad [\text{Newton's second law of motion}]$$

$$(2) \quad pV = nRT \quad [\text{The Ideal Gas Law}]$$

The central concept used for characterizing this transition is an idea a form of which we found in Brandom in §3.4, that of *explicit making*.

The rest chapter is divided into two parts. In the first part (§7.2) I note how explicit-making enlarges the toolbox for semantic functionalism, so that inference is only one of two kinds of fabric out of which a functional description can be built. Against this background, the picture described in §5 can be seen as providing a kind of identity semantic explanation for law statements. The key thought here is that law statements make explicit an aspect of a skill already mastered: the aspect of nomological connection between means and end. Earlier, in §3, I have noted the fact that semantic explanation of this sort cannot help with metaphysical debate, but rather presupposes that one party to the debate is right. Assuming that the diagnosis and therapy of §4-6 is success, this is now an acceptable presupposition. Characterizing our semantic practice in explicit-making terms gives something like a theory of experience, in an extended sense of “experience”, to which experimental success and mastery of skill also belong. This theory of (non-perceptual) experience complements the *epistemological* story developed in §5-6 the same way a theory of *perceptual* experience complements an epistemological theory of *perceptual warrant*. The theory of (non-perceptual) experience I arrive at has two components: first implicit information (“proto-content”) is taken up through the learning of a skill, then the information is, in an expressivist step, made explicit in an assertion.

The second part of this chapter (§7.3-4) deals with refinement of the picture developed in §5 to make it approximate the actual scientific practice more closely. What I aim at, is a *second* application of the idea of “explicit making” to describe the more sophisticated activities of modern science. Roughly, the key features of scientific investigation and experiments are efforts to “make explicit” various kinds of *responsiveness* in the mastery of simple, binary skills. These are primarily the implicit responsiveness to *degrees* and the implicit practical differentiation of *diverse factors*. These are aspects of binary skills that, unlike the aspect of nomological connection between means and end, were not made explicit in the picture developed in §5. Moreover, the latter kind of “explicit making” is no longer simply a matter of

Part IV: Beyond Therapy

“extracting” proto-contents in non-representational skills and putting them into explicit, discursive forms. Numerical measurements, formulation of mathematical functional relationships, as well as separation of and control for factors are all *further developments of skills*.

The description of a type of claims as “explicit making” presupposes the intelligibility of the contents expressed by these claims. The right to the above description of complex scientific law claims as making explicit various kinds of practical responsiveness implicit in binary skills, therefore, must be earned by showing that these scientific law claims enjoy a presumption of contentfulness, the way binary law claims enjoy a presumption of contentfulness. The way to do that is to show that the skill-expansion just described provides *warrant* for complex scientific laws in a manner similar to – though also with significant contrasts – the way binary skills provide warrant for binary law statements. Having secured the warrant for scientific laws we can finally cash the promissory note: to describe the process of skill expansion as making explicit nomological information.

Explicit making is however not the only type of development of complex and linguistically loaded skills that make up scientific activities. Another type of skill-expansion crucial for science is the kind of conceptual innovation known traditionally as “postulation of theoretical entities”. I shall indicate briefly why such empirical conceptual innovation of science does not threaten the broad realism defended in this essay.

§7.2 Law Statements Make Explicit Skills

§7.2.1 The Dialectical Division of Labor

The stated purpose of §5 was to show that, when we correctly describe the rudimentary elements of the nomological practice, the supposed disharmonies of language use with which anti-realists have motivated a skeptical dialectic simply do not appear. The goal was therapeutic, to the extent that it aims at *stopping* a philosophical urge to become a nomological anti-realist, rather than at *constructing* a theory. But the picture sketched in §5 also provides the basis for a positive *identity semantic explanation* for law statements. That is, it provides the basis for an answer to

the question: what makes law statements have the *kind* of meanings they have (i.e. nomological ones) ?

It is important to understand the dialectical division of labor here. The therapeutic part of the project is to secure the presumption of naïve realism about laws by undercutting the crucial presupposition of the skeptical dialectic – that of disharmony of use – thereby shifting the burden of argument back to the anti-realist. Success of the therapeutic part means that we are entitled to presume the coherence of law statements and the intelligibility of the kind of meanings they express. The second, *explanatory* part of the project, *assumes* that the therapy of §5-6 was successful, and that consequently, laws are real and nomological meanings are intelligible. What is being explained, is therefore *not* how nomological meanings are *possible*, but what makes law statements express these meanings, rather than some other kinds of meanings. To proceed, we need to make an excursion on some general features of explanation and their manifestations in semantic explanations.

§7.2.2 Interlude – The Context Sensitivity of Semantic Explanations

§7.2.2.1 Van Fraassen & The Pragmatics Of Explanation

Phillip Kitcher and Wesley Salmon wrote that van Fraassen has, with the chapter “The Pragmatics of Explanation” in his book *The Scientific Image*, delivered the best theory of the pragmatics of explanation, but that the theory “faces serious difficulties” if it is seen as a pragmatic theory of explanation.²⁷⁷ This is in my view an accurate assessment. In the rest of this chapter, I shall spell out, à la van Fraassen’s theory about the pragmatics of explanation, the extreme, but routinely neglected, context sensitivity of semantic explanations, especially that of identity explanations.

Van Fraassen thinks that there is a one to one correspondence between “why”-questions and explanations. An explanation is an answer to a “why”-question, while a “why”-question is a request for explanation. What counts as a satisfactory answer to a “why”-question depends on the context in which the question is put or entertained. Van Fraassen distinguishes two sorts of context dependence of “why”-questions, and hence of explanations. The first is what he calls the “interests” of those who discuss the question. He mentions the Norwood Russell Hanson’s example to illustrate this

²⁷⁷ See Kitcher & Wesley’s Van Fraassen On Explanation.

Part IV: Beyond Therapy

sort of contextual variation.²⁷⁸ In Hanson's example, the same death is explained by a physician in terms of "multiple hemorrhage", by a lawyer in terms of "driver's negligence", by a carriage builder in terms of "a defect in the brakeblock construction" etc. The second sort of context dependence is what van Fraassen calls the dependence on the "contrast-class" determined by the context. The idea of a contrast class is the idea that, in asking "why *p*", one is implicitly contrasting "*p*" with some other states of affairs that did not obtain. And with which the contrast is made, can vary from context to context. To use van Fraassen's example, in asking "Why did Adam eat the apple?" one can be contrasting the fact that Adam ate the apple with possible states of affairs where *someone else* ate the apple, or with states of affairs where Adam ate *something else*, or with states of affairs where Adam *did something else* to the apple. Each of these contrasting class gives rise to a different question, and therefore also a different potential explanation.²⁷⁹

§7.2.2.2 Semantic Explanations

Like all explanations, the identity semantic explanation and the *ti-esti* semantic explanation both have implicit context dependencies. To see what these might be, let me put the corresponding requests for explanation in the form of "why"-questions:

(Identity "Why"-Question): Why does *E* mean what it means?

(*Ti-Esti* "Why"-Question): Why does *E* have a (linguistic) meaning at all?

What are the contrast-classes for each of these questions? Now I think it is clear that, when a philosopher asks the *ti-esti* question, say about the expression 'red', he is not contrasting it with the situation where 'redd' would have meaning. For if the question is interpreted as "Why does 'red', *but not 'redd'*, have a (linguistic) meaning?", the answer might be: because early English writers decided to adopt the spelling 'red', and to reject the spelling 'redd'. So interpreted, the question and answer deal in historical philology, rather than in philosophy. The intended contrast class is hinted at by the phrase "at all" at the end of question, which points to the interpretation "Why

²⁷⁸ C.f. Van Fraassen's *The Scientific Image*, page 125ff.

²⁷⁹ The contrasting class can often be expressed using various means of *focus*, for example, certain accentuation in the articulation of the why-question. For a good survey topic on focus and information-structure expressed by it, see Rooth's "Focus".

does ‘red’ have a meaning, *rather than being meaningless?*” But fixing the contrast class this way does not free us from the threat of historical philology. For the answer to the latest question may still be “because early English writers decided to adopt the spelling ‘red’ for characterizing the color red; if they had chosen some other spelling, ‘red’ would have been meaningless”. To escape historical philology, we have to fix the “interest” of the interlocutors. In the present case, we are interested in the “constitution” of meaningfulness, rather than historical accidents about which spellings were adopted and which not.

In fact, my original formulation of the requests for explanation already eliminates such deviant interests as historical philology:

(Identity Constitutional Question):

In virtue of what does *E* mean what it means?

(*Ti-Esti* Question): In virtue of what does *E* have a (linguistic) meaning at all?

The constitutional formulation with the W-phrase “In virtue of what” not only fixes the interest, it also seems to eliminate choices of contrast classes. But this is not so. We have seen that we can think of the identity question, constitutionally formulated, in terms of the *ti-esti* question *plus* a differential, which is the non-constitutionally formulated identity question.²⁸⁰ We have seen that the *ti-esti* question, constitutionally interpreted, implicates difficult issues about normativity and the mind, and will not be taken up in the rest of this essay. It is the non-constitutional identity question that I will be interested in answering, in connection with law statements.

The identity question is highly context sensitive. Again, we need to fix the “interest” of question-mongers to escape historical philology. For we are not interested in answers like “because the spelling ‘red’ was adopted at such and such a time to describe things as having the color red, rather than as an interjection”. But the identity question is particularly sensitive to contrast-classes, and there is a great range of variation for the choice of a contrast-class that is capable of yielding a legitimate

²⁸⁰ See page 88 above for a precise formulation of the decomposition of the identity constitution question.

Part IV: Beyond Therapy

question. The intended range is somewhat narrowed by the fact that I have named the question “*identity question*”: it was precisely to prime us for interpreting the question as “Why does ‘red’ mean *red* as opposed to something else?”. But many different things can go in for “something else”. To appreciate the degree of variation possible, let us start from the other end, as it were, and look at the following possible *answers* to the identity question about ‘red’:

- (a) ‘red’ is used to classify objects;
- (b) ‘red’ can be used to make perceptual judgments/observation reports;
- (c) ‘red’ can be used to make perceptual judgments/observation reports *of a certain kind* (e.g., those that are based on/warranted by perceptual experiences characterizable as „seeing an object to be red“.)

Of these possible answers, it might appear that (a) is a fake. But this is not so. (a) is a perfectly good answer to the identity question, if the contrast class is “as opposed to the kind of meaning proper names or sentences have?” But if the contrast class is “as opposed to the kind of meaning semantic vocabulary – such as ‘true’ – has?”, the answer (a) ceases to be a satisfactory answer. For ‘true’ is also, on the face of it at least, used to classify objects (namely sentences). For that contrast class, (b) is a good answer (‘true’ is not used to make observation reports). With yet other contrast classes, (c) replaces (b) as an adequate reply, these are “as opposed to *hot*, or *blue*, etc.”.

Having narrowed down our interest to (non-constitutional) identity explanation, therefore, is not yet to have reached a definitive form of the kind of semantic explanation I want to pursue. Exactly which contrast-class will be used in the identity semantic explanation I develop for law statements, consequently, will be an issue of some interest later.

§7.2.2.3 Identity Explanation & Functional Roles

The examples I gave above for answering the identity question about ‘red’ might have surprised the reader. For these answers do not appeal to the *functional role* of the word ‘red’ in any genuine sense of “functional”. Those answers were more *relational* as functional.

It is true that the kind of answer that an *inferentialist* would give to any identity questions about meaning would have the form: in virtue of the inferential roles the expression in question plays in our language practice. The basic example from which inferentialists tend to generalize is that of logical constants. Logical constants remain arguably the most convincing case for this type of strict inferentialist answers. We saw in §4 that in *The Logical Basis of Metaphysics* Dummett sought to do something along the *relationalist* lines: he tried to establish justificatory procedures with which the entire set of use-rules of a logical constant can be “derived” from its introduction-rules (or its elimination-rules) alone. If successful, this would give rise to identity semantic explanations of logical constants by reference to their introduction-rules *alone*. Such explanations would have a simple relational form, relating a logical constant to the premises form which it can be introduced. There are, as we noted earlier, some somewhat technical reasons, why such a relational account of the meanings of logical constants would not work.²⁸¹

Whatever the case is with logical constants, there is reason to think that, at least in some other cases, appeal to the *full* inferential role is not required for an identical semantic explanation. So the mastery of *some* intra-linguistic inference rules is probably necessary in order to count as *understanding* the term ‘red’, and the word’s being used in accordance to these rules will probably have to enter into an account of what makes the word *meaningful* at all. But to explain what makes ‘red’ mean *red* as opposed to *expensive*, or *true*, or the person Napoleon, or the proposition that *It is too early to tell whether the French Revolution was successful*, it is probably sufficient to point out that the word is used to differentially respond to red things. This is the case for most contrast classes. It takes a very special contrast class for it to be necessary to mention intra-linguistic inference rules. An example of that might be “what makes ‘red’ mean *red* as opposed to ‘reflecting in daylight light waves with wavelength in the range R’?” (contrast class underlined), where ‘R’ is substituted for by an appropriate numerical specification. Since the complex physical description underlined can also be used in differential responses to red things, it appears necessary, in order to answer the question with *this* contrast class, to appeal to, for

²⁸¹ For a brief discussion of these, see the section on “Total Harmony or Intrinsic Harmony”, §4.1.4, below.

example, intra-linguistic *inferential* connections ‘red’ has to ‘colored’, but not to ‘emitting/reflecting photons in daylight’.

As it turns out, the kind of contrast classes most interesting for the identity question about law statements do not require a genuinely *functional* explanation. To this I now turn.

§7.2.3 Identity Semantic Explanation of Law Statements

We have seen that identity semantic explanations are highly context sensitive, and that they are especially sensitive to the contrast-class implicitly understood with the question. What is the implicit contrast class in the question “in virtue of what do law statements have the meanings they have?”? Some contrast classes are utterly uninteresting, for example “as opposed to the kind of meaning a singular term has”, or “as opposed to kind of the meaning a predicate has”. To be interesting, we have to contrast the meanings of law statements with other kinds of *propositional meanings*. Given the way the bulk of the philosophical disputes about laws of nature have been going, it is quite obvious that the most interesting propositional meaning contrast is with meanings of (i) singular observation reports, and (ii) corresponding universally quantified statements. In other words, the most interesting identity semantic question (IS-question for short) for law statements has the following form:

IS-Question for Law Statements: What makes a law statement have the meaning it has (i.e. a nomological meaning), *as opposed to* (i) the kind of meanings observation reports have, or (ii) the meaning of the corresponding universally quantified statement?

The picture of §5, as primitive as it is, contains enough material to answer *this* question. As it is the case with the identity semantic question for empirical predicates²⁸² (for most contrast classes), no reference to the *entire range* of intra-linguistic usage is necessary. So one answer to IS-question for law statements, appealing to the kind of *warrant* available to law statements, might be: in virtue of the fact that the mastery of a certain skill, rather than a perceptual experience, or several perceptual experiences, can provide direct warrant for a law statement. By contrast, an

²⁸² I briefly explained the reasons for thinking this in §2.5.3 above. Detailed argument for this position will lead us too far afield.

observation report can be directly warranted by a perceptual experience of the right kind. A universally quantified statement can be directly warranted either by (a) several observation reports of the form “*this* A is B” and what I call an *enumeration claim* of the form “these are *all* the As there are”, or by (b) a proof from “N is A” to “N is B”, where the basis of this proof can be either some conceptual truths, analytical truth, or even laws of nature. Either way, a universally quantified statement is *not* directly warranted by the mastery of certain skills.

If the mentioning of direct warrant to explain the contrast with universally quantified claims ((ii) in the statement of the IS-question above) does not bring out clearly the relative *strength* of law statements vis-à-vis universally quantified claims, there is another explanation available: a law claim has the meaning it has as opposed to the meaning of a corresponding universally quantified claim *in virtue of the fact that it, but not the universally quantified claim, can license certain kind of desiderative inferences in practical deliberation*. That is, the explanation for the meaning difference between law statements and universally quantified statements can point to the kind of *consequences* each is capable of supporting.

§7.2.4 Explicit Making As A Fabric Of Functional Explanation

Pointing to the warrant and consequence of law statements – the “verification” side and the “pragmatic” side of use, respectively²⁸³ – can therefore yield answers to the IS-question about law statements. I would like to provide another kind of answer, based on an idea from Brandom we have examined earlier, the idea of “explicit making”.

There are two reasons for adopting this kind of IS-explanation. The first is that it enables me to achieve *explanatory continuity* with the story I am about to develop for more sophisticated laws statements that appear in the context of *epistémé*: the natural sciences. For as it turns out, much of the sophistication in the sciences consists in their more *in depth* explicit making, a form of which is already taking place in the relatively primitive practice we have sketched in §5. The second, more fundamental reason is that explicit-making gives a more refined picture of our epistemological access to laws of nature, through a kind of theory of (non-perceptual) experience.

²⁸³ For more on this Dummettian terminology, see §4.1.4 above.

Part IV: Beyond Therapy

Instead of saying, simply, *that* (binary) skills provide warrant for (binary) law statements, we can further spell out *how* this warrant works. The idea is that, in acquiring (binary) skills, we take up, *non-representationally*, nomological information, or *proto-content*. This proto-content is then made explicit in an expressivist act of asserting the corresponding law statements.

One formulation of the alternative, explicit-making based explanation is this:

(Nomological Explicit Making) A (binary) law statement has the meaning it has, because it is used to make explicit a nomological *proto-content* contained in a (binary) skill a subject (or a community) has mastered.

The Nomological Explicit Making thesis has the form of what I have called *vertical content explanation*.²⁸⁴ So it assumes the intelligibility of nomological information. Now we are entitled to that assumption because of Part II & III. As noted, the explanation further specifies the underlying *warranting relation* between mastery of a skill and the law statement based on it. We saw in §5 that the relation is not like the warranting relation between a perceptual experience and the observation report based on it.²⁸⁵ But nor is the relation like that of *inferential* warrant. The mastery of a skill, or an exercise of that mastery for that matter, is not a *representation*, let alone an assertion. Hence the warrant provided by mastery of a skill is unlike the warrant provided by the premises of an argument for its conclusion. The Nomological Explicit Making thesis claims that the warrant in question has the form of a *proto-content* bearer giving warrant to a statement that explicitly *asserts* that *proto-content*. This relation is more like the warrant provided by a good inference pattern – which bears a certain “argumentative connection” as a *proto-content*, though does not *assert* it – to a corresponding conditional claim, *provided* it makes sense to speak of an inference pattern as bearing *proto-content*. But skill-warrant also differs from warrant by good argument in that the *proto-content* in the mastery of a skill is *empirical, nomological* information.²⁸⁶

²⁸⁴ For the contrast between vertical and horizontal content explanation, see §3.4.2.2.

²⁸⁵ C.f. §5.3.3 for more details on this claim.

²⁸⁶ I do not claim that good arguments *do* have proto-information. I am using the hypothetical case to compare and contrast with the case I *do* want to make a claim about.

The relation between the thesis that skills provide warrant for law statements and the Nomological Explicit Making Thesis is analogous to the relation between (a) saying that perceptual experience provides warrant for an observation report and (b) giving a *theory of perceptual experience*. I have insisted, in §6.2.1, that experimental success and skill mastery is also a form of experience, no less than *perceptual* experiences are. The Nomological Explicit Making Thesis says that skill mastery is part of a larger experience, the other part being the explicit-making of the nomological information implicit in the skill mastered. As a matter of fact, I think a plausible theory of experience *also* should appeal to the notion of explicit-making. But that is a topic for another occasion.²⁸⁷

Before moving on, I need to say more about the notion of *proto-content* itself.

§7.2.4.1 *The Notion Of Proto-Content*

The idea of *proto-content* I would like to make use of in semantic explanation has five components:

- (PC1) Proto-content is information;
- (PC2) Proto-content is not a form of *representational* content, in the sense that its bearer is not a representation;
- (PC3) Bearers of proto-contents are *conscious* episodes or capacities for such episodes on the part of a subject;
- (PC4) Proto-content can potentially be made explicit – through conscious reflections and acts – by its subject in a claim; and it can, but need not be identical with the content of that claim;
- (PC5) A proto-content and the corresponding explicit-making claim have the same subject-matter, and it is not the *bearer of proto-content*.

Let me take up (PC2) and (PC3) first. Bearers of proto-contents are conscious, but non-representational episodes (including activities) or capacities for such episodes/activities. Examples of such bearers are: perceptual experiences, the drawing

²⁸⁷ C.f. my remarks at the end of §3.4.2.3 above (especially the footnote there), as well as footnote 288, and the paragraph on page 292 about making explicit perceptual experience below for a bit more detail.

Part IV: Beyond Therapy

of an inference, skillful manipulations. Of these three examples, the first is controversial and I will not go into it.²⁸⁸ The third is the present topic to which I shall come presently. The example of drawing an inference can serve to illustrate the idea. One is not only *conscious* when drawing an inference. In fact, the drawing of inference *includes* parts that *are* representational or at least *means* for representation (e.g. the utterance of the premise and conclusion). Yet still, to draw an inference is not to make a claim, or to represent the world in some other way.

The requirement of non-representationality is to ensure that we are talking *not* about claim-contents, but about *proto*-contents. The requirement of consciousness is necessary for (PC4): the possibility of being made explicit. If I am not aware in²⁸⁹ the episodes carrying a proto-content, I cannot possibly make a statement based on the proto-content bearer. (PC4) also says that to be made explicit *can* mean being identical with the content of a claim. Here I am assuming a genus of informational contents, of which proto-contents and the contents of claims are species. Informational contents are specified by sentences, but not all bearers of informational contents are claims, or even linguistic. They are not when the informational content is merely *implicit* in Brandom's sense. On the other hand, the content of every claim is also an informational content, for content asserted is also information delivered. So the identity envisaged here is between the proto-content and the content of the explicit-making claim *as* informational contents.

For some cases, specifically the case of perceptual experience, I do not think the proto-content in an experience is *identical* with the content of the observation report based on it. The reason is that a judgment has a *much more determinate* content than the information carried by a perceptual experience. Similarly, in the case to be discussed shortly, the information carried by my degree- and factor-responsiveness that is implicit in my mastery of a skill is also less determinate than the numerically expressed measurements and equations that make this responsiveness explicit. But in

²⁸⁸ Most philosophers think that perceptual experiences *are* representational; I think they are representational only in a derivative sense, relative to the possibility of forming perceptual judgments based on them. In themselves, perceptual experiences are as representational as an accidental imprint. This position presupposes some notion of *non-conceptual* sense impressions, such as developed in the first chapter of Sellars *Science and Metaphysics*.

²⁸⁹ Crucially, not "aware of": I do not need to think *about* my perceptual experience to make an observation report on it.

both these cases, the implicit and explicit informational contents are not unrelated. I shall say more about the case of responsiveness in skills later. The point here is to illustrate how the claim-content can come apart from the proto-content, not by being unrelated to the latter, but by being more determinate than it.

(PC5) says that the explicit-making claim is not *about* the bearer of the implicit *proto-content*,²⁹⁰ rather, it has the same “subject matter” as the proto-content. So, an observation report is not *about* the experience it is based on; a conditional is not *about* the drawing of inferences,²⁹¹ and similarly, a (binary) law claim is not *about* a (binary) skill or its exercises.²⁹² The idea can be paraphrased this way: what differentiates the explicit claim and the implicit proto-content bearer is the *addition of intentionality* in the explicit claim; they need not differ in informational content. A claim purports to *represent*, a visual impression or the drawing of an inference do not. *What* the explicit claim represents is not the proto-content bearer. It is rather what the proto-content bearer (e.g. the endorsing of an inference) *would have* represented, if the latter *had been* representational (which of course it *could not* be).

The first requirement (C1) might raise some brows, especially if I am going to use explicit-making to give an account of nomological contents. For it might be thought that the concept of *information* must be explained by modal concepts, such as counterfactual robustness. Another worry is that, to speak about *nomological information*, one would have to have a notion of *second-level* nomological connections, in order, for example, to conceive of how such information is *transmitted*. The way to respond to these worries is first to remind ourselves that the modest form of the semantic explanation here *presupposes* the intelligibility of nomological and modal contents, and aims to say what makes law statements latch onto nomological meanings. Moreover, transmission of nomological information

²⁹⁰ Unless the proto-content is about its own bearer. This sort of self-referentiality is not a realistic possibility for our purpose.

²⁹¹ But about certain “argumentative connection”, let us say.

²⁹² As a side remark, this is how we must make sense of Brandom’s claim that conditionals make explicit endorsements of inferences. The content of a conditional – insofar as it makes sense to attribute contents to conditionals – is not the endorsing *act*, nor the endorsing *attitude* on the part of the subject, nor is it deontological, about the *proprieties* of so endorsing. Rather, it has the same subject matter as the *proto-content* carried in that endorsing act or attitude, which must be understood in a more objective manner, as some sort of “argumentative connection”.

Part IV: Beyond Therapy

through, say the acquisition of a capacity, need not require anything like a conception of second-order nomological connection to which the defense of §5-6 did not (even aim to) secure. This can be illustrated by the fact that we can explain, using first-order scientific laws, how a piece of iron acquires the *disposition* of elasticity (the explanation involves a description of how the atoms and subatomic particles in the iron bar come to be aligned in a certain way). To conceive the transmission of nomological information, similarly, should not require the concept of second-level nomological connections. Entitlement to first order nomological connections that it *does* presuppose, is secured by Part II & III.

§7.2.4.2 Nomological Connection Between Means and End

It is easy to see that (PC2-3) are satisfied by (binary) skills: a skill is a capacity for a certain kind of conscious, non-representational actions. My claim now is that we can think a binary skill as carrying information that is, through a conscious reflection, made explicit into a corresponding law statement. This information, as content, is identical with the content of the binary law statement.

Recall that mastery of skill may presuppose the ability to *formulate in words* one's goal and the means for achieving it. What is made explicit in the binary law statement is not a type of goal, nor a type of means. For these are already explicitly represented by those who master the skill. What the law makes further explicit, is the *nomological connection* between a type of means and a type of end. For this reason, we could also think of the *proto*-content made explicit by a binary law statement as a "nomological connection between means and end". This is however only *one* aspect of the binary skill made explicit. Next we shall see that attaching a *ceteris paribus* clause to a binary law statement, or the development of scientific laws, are other ways of making explicit other aspects of the mastery of binary skills.

The theory of (non-perceptual) experience complementing the epistemological story of §5-6 which we have arrived at, then, is this. First, by acquiring a binary skill – for making an S-situation B by making it A – we take up, in non-representational form, an implicit *proto*-content storable as a nomological connection between the means and end for which the skill is a skill. Then, in a reflective and expressive act, we make explicit this nomological connection implicit in our mastery of the skill by asserting the corresponding binary law claim.

§7.3 Implicit Mastery of Techné

Recall how we characterized binary laws of the form (LS)²⁹³ (“T” for *Techné*):

- (Ta) strongly qualified by *ceteris paribus* clauses;
- (Tb) simple: has only two terms, ‘A’ and ‘B’;
- (Tc) isolated, rather than belonging to a system, inter-connected through, say “derivation” relations;
- (Td) no essential use of mathematics, nor other kinds of numerical device.

Consider for example the following binary law statements:

- (Match Law) It is a law of nature that if a match is struck, it will be lit.
- (Gas Law) It is a law of nature that if more gas is pressed into a chamber, its pressure will increase.

Both laws have the generic form (LS) discussed in §5.1.3, and are “binary” in that they “link” two terms ‘A’ and ‘B’ together, though each term can have complex contents. This is character (Tb). They are also not true, strictly speaking, unless understood with a *ceteris paribus* clause. (Character (Ta)). A wet match, or a gas chamber whose walls are connected with a cooling system, respectively, in addition to countless other circumstances, will defeat these laws. Moreover, these two laws are, because (as we have been supposing) connected with individual skills, not apparently connected with other law statements in any thing like a system. (Character (Tc)). Finally, none of these laws uses mathematics. In fact, since the skills underlying these laws include no *measurement skills*, we are to think of ‘pressure increase’, or ‘compressed’ as primitive *ungradable* predicates. (Character (Td)).

In contrast, both Newton’s Second Law of Motion and the Ideal Gas Law shown above have the following, contrasting characters (“E” for *Epistemé*):

- (Ea) The law statements are, at least on the surface, *absolute*.

²⁹³ C.f. §5.1.2 above.

Part IV: Beyond Therapy

- (Eb) They can be highly complex, involving many terms in complex mathematical relations;
- (Ec) They are systematically related to other law statements, through “derivation” relations;
- (Ed) They involve extensive use of mathematics and numerical methods.

The question is, given a practice described in §5, what sorts of additional practice enables us to make knowledge claims about *scientific* laws: laws that exhibit the characters (Ea)-(Ed)?

§7.3.1 Implicit Mastery

The first step is to recall that, binary skills as described in §5, as well as corresponding desiderative inferences in practical deliberation, have an essentially *implicit* aspect.²⁹⁴ What I want to suggest now, is that we can distinguish two components in this implicit aspect. The mastery a skill, as well as its exercise requires, to varying extent, the ability to *be responsive to degrees*, as well as to *respond differentially to various factors*.

§7.3.1.1 Implicit Responsiveness to Degrees

To start, consider the (Match-Law) example. A person who has the corresponding skill must also know, in a *practical sense*, how quickly to strike a match, for how long to keep the match in contact with the strike-surface, how firmly to press the match, and perhaps also how rugged and hard the surface he should strike the match against. Although a person capable of lighting a match by striking it need not know any mathematics, nor be able to formulate precise measurements of any of the dimensions just mentioned, he must be *responsive* to these dimensions to be said to mastered the skill.

The same is true of the person who is capable of increasing the pressure of a chamber by putting more gas into it. The example we have chosen may look somewhat artificial, but perhaps inevitably so. For we want to consider a primitive,

²⁹⁴ See §5.2.2.2 (implicit understanding for *desiderative inferences*), §5.3.1 (implicit mastery for *skills*), and §5.4.1 (implicit harmony between the two aspects).

binary form of a scientific law – in this case the Ideal Gas Law – in order to observe, in a thought laboratory as it were, what is needed for it to metamorphosize into its scientific form. The binary skill – to increase the pressure of a chamber by putting more gas into it – appears to be an artificial one in that it seems to be not a useful skill as the skill for lighting a match by striking it is. Yet since the way from the latter useful skill to scientific laws is much too long and serpentine, it is not a good illustration after all.

As it turns out though, there *is* a context in which the pressure-gas skill is useful: to increase the pressure of a (bike) tire by pumping air into it. Now, like any useful skills, this skill is robust under variations of degrees. That is, it cannot be a skill for increasing tire pressure by or to some *particular extent* (say 5%). It must be a skill for increasing tire pressure for a range of values that may be desired (though of course not explicitly in numerical forms). We learn quickly to judge, independent of the type of tires and the kind of air pump we use, how much air is needed in order to give the tires of a bike enough pressure for a pleasant ride. And we can learn that without ever learning to read a manometer, or counting the total number of times we have to press the air pumps. In other words, then, the skill in question incorporates an implicit responsiveness to the relation between the amount of air and the pressure achieved.

§7.3.1.2 Implicit Differential Responsiveness to Different Factors

The other component of the implicit aspect of the master of a (binary) skill is the implicit responsiveness to the practical significance of various factors. In the case of lighting a match by striking it, a person cannot be said to have acquired the skill if he is not responsive to whether the match is dry in his attempt to light it by striking. The dryness of the match, like the windiness of the location, is a factor the subject must be responsive to in order to count as capable of lighting a match by striking it. The point here is that the subject must be sensitive to a *variety* of factors, regardless whether these factors are *gradable* (but in case some are, the required sensitivity must involve a sensitivity to the *degrees* of these factors, as illustrated in the previous paragraph).

The situation with the pressure-gas skill is no different. We all know that the pressure of the tires can change when *temperature* changes. So if we are pumping air into a bike tire for a ride outside, in the depth of wintry Siberia, we will freely put a bit more air into the tires. But if we are charging the tires in the early morning for a

Part IV: Beyond Therapy

daylong bike tour in the height of Sicilian summer, we will make sure to put a bit less air into the tires as we normally we would. In short, the pressure-gas skill requires sensitivity to *ambient temperature change*. Although the factor of temperature change is not *explicitly expressed* in the way we describe the skill: to increase the pressure of a (tire) chamber by putting more (air) gas into it. The *actual mastery* of the skill must include the development of a sensitivity to this implicit factor.

§7.3.2 *Ceteris Paribus* Clause of Binary Laws

The importance of the implicit aspect of a binary skill – especially the two components of implicit responsiveness just discussed – for the corresponding *law* statement, is that it underlines the fact binary law statements are felt to come with *ceteris paribus* clauses. One can show that the (Match Law), for example, does not hold strictly, by saying “supposing the match is wet...” or “supposing that it is extremely windy...”. Such objections suggest that the law statement is only true when the clause “*everything else being equal*” is attached to it. The objections to unqualified truth just cited spell out, as we might say, two *defeating circumstances* for the unqualified (Match Law). What is interesting for us is that these defeating circumstances are correlated with the factors to which *practical responsiveness* is indispensable for skill mastery. The correlation can be thought this way:

a *failure* to be adequately responsive to one relevant factor in the skill creates a defeating circumstance for the corresponding law statement.

The other component of the implicit aspect of skills – practical responsiveness to *degrees* – is also correlated with defeating circumstances this way. Consider the (Gas Law). One defeating circumstance of the law is when the “more gas” is a tiny amount, not enough to offset the sharp drop of ambient temperature. (If I just put a tiny amount of air into the tires of my bike while in the warmth of my home, and then take the bike outdoors in extreme winter, the pressure of the tire will likely *decrease*, not increase.)

There is much debate about the logic and function of such clauses. It is very likely the phrase does not have a uniform meaning or function in all contexts in which it appears. But one might quite generally characterize the function of *ceteris paribus* clauses as one of *signaling the existence of defeating circumstances*. They might do

more, but they at least do this much. Note that that is a *truth-conditional* way of thinking about *ceteris paribus* clauses, where “defeating circumstances” are just the complement, as it were, of “truth conditions”.

What I want to suggest now is that, for *ceteris paribus* clauses attachable to *binary law statements*, it is more helpful to think about these clauses not truth-conditionally, in terms of defeating circumstances of the *law claim*, but *expressively*, in terms of making explicit a *generic* piece of information implicit in the *skill-mastery* that provides warrant for that claim. More precisely, and using our term of art, “*proto-content*”, I claim that:

(Expressivist Thesis About *Ceteris Paribus* Clauses) A *ceteris paribus* clause “everything else being equal”, when attached to a binary law statement, makes explicit a certain fact about the conditional nature of the binary law, a fact that is, as *proto-content*, implicit in the *practical responsiveness* of the mastery of that binary skill which warrants the law statement.

The idea needs some spelling out. We have seen that the very necessity for skill-mastery of *some* practical responsiveness of either of the two kinds just discussed is correlated with the *existence* of defeating circumstances for the unqualified binary law statements. The necessity for practical responsiveness – regardless *which* kind and to *what* – can be thought of as a feature of the skill that carries the *information* that, if we were to put it in words, can be expressed by saying “there are defeating circumstances for this nomological connection”. Instead of thinking of the *ceteris paribus* clause as signaling the possibility of defeating circumstances, therefore, we can think of it as making explicit this bit information carried by the *unspecific* feature of the skill: its mastery requires *some* practical responsiveness.

It is important to understand what *kind* of claim the Expressivist Thesis about *Ceteris Paribus* Clauses is. It does not give the *logical form* of the clause, nor does it analyze or even paraphrase the clause in other terms. What it purports to do is to give a broadly functionalist semantic explanation to the following identity semantic question about binary law statements:

Part IV: Beyond Therapy

(*IS Question about Attacheability of Ceteris Paribus Clauses*) What makes it the case that a binary law statement has this semantic feature: what it says can be said with a *ceteris paribus* clause attached to it?

The answer provided by the Expressivist Thesis, is that (i) a binary law statement makes explicit the nomological connection between means and end that is implicit in the mastery of the binary skill that warrants the law statement, that (ii) part of the mastery of this binary skill is some (unspecified) practical responsiveness (to degrees and factors), and that (iii) attaching a *ceteris paribus* clause makes explicit, beyond the mere nomological connection between means and end, what is implicit in *this* aspect (i.e. (ii) above) of the skill-mastery.

In making explicit the practical responsiveness, the *ceteris paribus* clause attached to a binary law statement does not distinguish the various *proto*-contents carried by *each specific responsiveness* implicit in the mastery of the corresponding skill. To make explicit the *proto*-content in these specific responsivenesses *individually*, is to take the first step towards science.

§7.4 From Techné to Epistémé

§7.4.1 Justifying vs. Explaining Scientific Laws

One thing is clear: simple binary skills *cannot* provide warrant for precise, mathematically formulated, and systematically correlated scientific laws. Unlike the attachment of *ceteris paribus* clauses to binary law statements, for scientific laws, new warrant has to be *earned*. Part of this earning process is to expand binary skills.

Provisionally, this process of skill expansion can also be thought of as a process of making explicit the *specific* responsivenesses to factors and their degrees in already mastered binary skills. Take the (Gas Law) as an example. We saw that mastering the underlying skill requires a practical responsiveness to ambient temperature changes. This practical responsiveness can now be *verbalized* in articulating the conditions of exercising the skill: “when the ambient temperature stays constant”, for example. The other practical responsiveness we mentioned – to the correlation of degrees between pressure and amount of gas added – can be replaced by explicit *measurement claims* as well as *mathematical claims of functional*

dependence.²⁹⁵ The idea here is *not* that the precise measuring claims and mathematical functional relationships are *already* implicit in the mastery of the binary skill for increasing pressure by pumping in gas. Measurement claims and mathematically formulated claims are, to the extent they are representational, express *more determinate* information as the practical responsiveness in the binary skill carries.

In any case, the picture of explicit making is provisional because we must *justify* treating measurement and mathematical relation claims as *representational*, and that requires laying to rest the question about the *warrant* for complex law claims. Semantic description and explanation, as we saw in the case of simple, binary law statements, can presuppose the success in answering the warrant question.

§7.4.2 Warrant

The first step towards warrant for complex law statements is to understand *measurements* and *exploration of mathematical relationships* as additional *skills* to be acquired. These additional skills are related to the old binary skill in the following ways:

- (A) The new skills, as capacities, *presuppose* the simpler skill;
- (B) The new skills, *taken collectively*, form one complex skill to manipulate the values of certain measurements;
- (C) The new skills, *taken individually*, are skills to make claims, and to that extent, are *explicit*;
- (D) The new skills, as individual skills to make claims, will express the same kind, but more determinate, information than the practical responsivenesses in the simpler skill carry;

For completeness, I have included in the list (C) and (D), which have to do, not with the warrant question, but with semantic explanation and description. There is in other words an explanatory caesura between (B) and (C).

²⁹⁵ I do not mean mathematical *laws*, but mathematically formulated claims about the relationship between various measurable quantities.

Part IV: Beyond Therapy

Let me then start with (A) and (B) and the warrant question. (A) expresses the simple idea that the ability to measure pressure and the amount of gas in the experimental setting of pumping gas to increase pressure obviously presupposes the skill to pump gas in order to increase pressure. (B) is more subtle. At first sight, the skill for measurement appears to have the form of a skill to make a particular kind of claims – *measurement* claims – in the style of, say, “the pressure of this chamber is...”, or “the ambient temperature is...”. Yet, the most fundamental way to understand these skills is not to think of them as individual skills to make claims. The picture is not that we acquire abilities to make new kinds of observation claims, each unrelated with the other. The picture is rather that, these skills are *components of one complex skill*, and this complex skill is not a skill to make *claims*, but a skill to *manipulate* pressured gas in a chamber, just *like* the simpler binary skill. The new complex skill is “merely” more flexible than the old one.

The picture at this stage therefore does not require thinking of mathematical and measurement claims as representational. Someone who has the ability to make pressure, temperature, and other related measurements as well as the ability to manipulate the following mathematical relation in conjunction with these measurements:

$$(2) \quad pV = nRT,^{296}$$

has a skill to bring a certain desired pressure about by controlling gas density and temperature. Just like the simple binary skill gives one warrant for asserting the binary law claim, mastery of the complex skill under discussion gives one warrant for asserting the complex law statement (2).

§7.4.3 Systematicity of Science

There is however a slight complication that did not exist for the case of binary laws. To be warranted to *claim* (2) presupposes that elements of the sentence express *concepts*. But each of the variables in (2) is associated with a certain *measurement skill*: the skill to make measurement claims. We cannot recognize (2) as a *claim* with

²⁹⁶ Where the variables are understood in the usual way: ‘*V*’ stands for volume of chamber, ‘*p*’ for pressure, ‘*T*’ for temperature, ‘*n*’ for amount of substance, and ‘*R*’ is a constant.

a representational purport without recognizing the various measurements as *concepts* having also a representational dimension.

Now for measurement activities to acquire a *representational dimension*, that is, for them to be thought of as measuring *some quantity in the world*, it must be possible to think of them as *independent of particular skills*. Consider the good advice that, when angry, count to 10 before saying anything. We might think of this as a skill for avoiding tempestuous outbreaks. One component of that skill is “counting to 10”. But there is hardly any sense that this counting, a measurement of some sort, is *representational* in the sense of reliably reflecting some quantity in the world. More likely the counting achieves the calming effect by *distracting my attention* away from the matter making me angry.

What this example illustrates is that measurement claims must achieve *autonomy* from particular skills, in order to acquire the status of making objective claims about quantities. What this means is that, in order for, say, the measurement procedures for arriving at the value for ‘*n*’ (substance amount) in (2) to acquire a representational dimension, they must also *reappear in other skills*, say the skills for manipulating chemical reactions.

This consideration confers a *holistic* aspect to the warrant for complex scientific laws. A person – more realistically, a community – does not obtain warrant for a complex scientific law by acquiring just *one* complex skill such as the one we described for (2). One has to acquire *many* such complex skills, so that the measurement *skills* involved in these complex skills become detached from particular complex skills, and thereby achieve the status of measurement *concepts*. Put it from the side of laws: one cannot be warranted to make the law claim (2) without simultaneously acquiring warrant to make law claims in which the *same measurements* – pressure, substance amount, volume, etc. – show up.

Notice that this holism does not include knowledge of *mathematics*. In order to credit (2) with the status of a claim about the world outside us, there is no similar presupposition that mathematical *laws* have to be representational about an aspect of external reality. What is required is merely that the mathematically *formulated* claims such as (2) are representational this way.

Part IV: Beyond Therapy

Finally, there is a second holistic dimension of scientific laws, due to the separation of heterogeneous factors. I shall briefly discuss it below in §7.4.5.

§7.4.4 Science as Expression

Once we are able to see that scientific laws have adequate warrant – a holistic one in the sense described – in the mastery of many complex skills, we are entitled to operate with the presumption of realism about the scientific claims.²⁹⁷ This presumption of scientific realism enables us then to give an a *semantic explanation* of why complex scientific laws such as (2) mean what they mean: they make (individually) explicit *proto*-information carried in the complex skills that give them warrant. Our talk of complex nomological *proto*-information, like our talk of simple nomological *proto*-information, presupposes the correctness of the presumption of nomological realism.

We are moreover in a position to give a description of the process of acquiring complex nomological skills on the basis of simple, binary skills. Recall what we said earlier in anticipation:

- (C) The new skills, *taken individually*, are skills to make claims, and to that extent, are *explicit*;
- (D) The new skills, as individual skills to make claims, will express the same kind, but more determinate, information than the practical responsivenesses in the simpler skill carry;

The skills to make measurements of pressure and the amount of gas, as well as the skill to predict, with the help of the mathematical relationship (2), measurement of the one given measurements of the others, can be thought of as refinements of various sorts of *practical responsiveness* to factors and degrees in the binary skill to increase pressure by adding gas.

²⁹⁷ An important element of scientific practice is yet to be introduced: conceptual innovation of theoretical entities.

§7.4.5 Absoluteness of Scientific Laws

By giving explicit expression to the factors to which one was only *implicitly* responsive to, *ceteris paribus* clause becomes less and less necessary.²⁹⁸ What gave rise to “defeating circumstances” are explicitly present as a variable in the formulation of a scientific law.

Sometimes, a factor is not taken up in the expansion of the original skill, as the factor of ambient temperature change in our example was. Sometimes, a factor is *controlled for and minimized*. A simple example is that skills associated with gravity minimize electro-magnetic effects. But because there is a “global” skill to compute the composition of gravity and electro-magnetic effects, *ceteris paribus* clause due to heterogeneous factors is not necessary, and is replaceable with an explicit instruction for the sort of “global” composition indicated. Here, in addition to the issue of measurement concepts discussed above, lies a second holistic dimension of science.²⁹⁹

The fact that we are able to view the progress from binary laws to scientific laws as a matter of *making explicit* various practical responsiveness implicit in the former, together with the fact that *ceteris paribus* clause is not felt to be attachable to scientific laws, further vindicates the Expressivist Thesis About Ceteris Paribus Clauses defended above.³⁰⁰

§7.4.6 Empirical Conceptual Innovations

I have, for the sake of clarity (or merely the impression of it as one might say), ignored a central practice of modern science: the creation of new empirical concepts, frequently known as postulation of “theoretical entities”. The elimination of *ceteris paribus* clauses, for example, is possible only with such conceptual innovations.

I will not have space in this essay for any detailed discussion of this important issue in the philosophy of science. The brief remarks I shall make are for the purpose of pointing at the direction one might go to maintain the broadly realist position – one

²⁹⁸ Conceptual innovation in the postulation of theoretical entities is needed to completely rid of *ceteris paribus* clauses.

²⁹⁹ C.f. §7.4.3 above for the first holistic dimension.

³⁰⁰ C.f. §7.3.2 *Ceteris Paribus* Clause of Binary Laws above.

Part IV: Beyond Therapy

that I have sought to strengthen in this essay by making it immune to Hume inspired skeptical dialectics – even in the face of such popular distinctions between theoretical and “normal” entities.

The creation of new empirical concepts such as that of electrons is fundamentally unlike the process of “explicit making” in terms of which I have described binary law claims and those scientific laws involving no such conceptual innovations. Inventions of such concepts can lead to wholesale changes in the practice of science, as famously argued by Kuhn in his classic *The Structure of Scientific Revolutions*, among others. There are broadly two things one might say in relation to conceptual (r)evolutions to defend the kind of nomological realism advocated in this essay. The first point is that conceptual (r)evolutions do not change the fact that we acquire, exercise, and attribute *skills*. The concepts of skill, practical deliberation, and experimental success, are all stable under such conceptual changes. Consequently, conceptual changes might put in doubt *which* nomological connections there are, and in *what form* they ought to be formulated. But they cannot put in doubt the *reality of skills*, and the reality of skills presupposes the nomological reality.

Moreover, the binary law claims are not affected by scientific conceptual change as much, because they are attached with such *ceteris paribus* clauses that signal an implicit aspect of skill mastery. How to make explicit this implicit aspect of a simple binary skill can be easily open to dispute, depending on the scientific conceptual paradigm one has in place. But *that* we have that skill, is more difficult to challenge, for it depends only on the goodness of *ordinary, non-scientific concepts*.

The second point to be made about scientific conceptual innovation and revolution is the Sellarsian point that perception of “ordinary” entities is not all that different from perceptions of “theoretical” entities. There are two sides to this point. On the one hand, empirical conceptual innovations involve not only *theoretical gimmicks* like bridge equations and ramsification procedures. They involve also the establishment of what Sellars calls *methodological correspondence rules*,³⁰¹ which establishes what counts as *observations* of a certain theoretical entities. On the other hand, even perceptions of “ordinary” entities require acquiring a complex conceptual

³⁰¹ For a good discussion of the contrast between substantive and methodological correspondence rules, see O’Shea’s *Wilfrid Sellars: Naturalism with a Normative Turn*, 27ff.

apparatus.³⁰² Science differs from ordinary perception in that it is a *self-conscious* attempt to develop a conceptual apparatus adequate to experience. But this difference ought not to make theoretical entities more vulnerable to anti-realism than “ordinary” entities.

³⁰² In addition to the usual “theory-laden-ness” of perception arguments, see also Sellars phenomenological discussion in his *The Role of Imagination In Kant’s Theory of Experience*.

Epilogue

Two central ideas explored in this essay revolve around skills. They are, first, mastery of skills as providing *warrant* for theoretical knowledge claims, and second, mastery of skills as a kind of *non-perceptual experience*. These ideas are tied together in the thought that, skills play a genuine epistemological role *by* being a kind experiential interaction with nature. The issue of anti-realism about laws of nature has served in this essay both to motivate and to test this nexus of ideas.

The problem of laws of nature itself was eventually approached therapeutically, and with the help of a third idea, taken from semantic functionalism: the idea of harmony. On the success of the diagnosis and therapy I shall leave the relevant chapters (§4-6) to speak for themselves. To close, it will be fitting to remark on the general philosophical significance of the ideas related to skills developed and applied in this essay. There are two aspects of that significance that I want to touch on. On the one hand, the view of skills developed in this essay differs from recent philosophy in the kind of roles they respectively make skills to play. On the other hand, this view of skills has non-trivial consequences for the empiricism that we know from Kant and its followers.

Skills and Epistemology

When Gilbert Ryle raised the importance of knowing-how versus knowing-that, he was not classifying two sorts of abilities. What Ryle meant is that knowing-that *is* a form of knowing-how. The role of skills for Ryle is one of explaining the “intelligent character” of certain human behavior without appealing to private mental episodes. Intelligence, as well as knowing-that, is to be *explained* by appeal to knowing-how. No epistemological relation between knowing-that and knowing-how was envisaged as part of the picture.

What is true of Ryle is true of a diverse group of philosophers who have emphasized skills. The three figures discussed in this essay in connection with semantic functionalism – Sellars, Dummett, as well as Brandom – all have a tendency

to *explain* linguistic understanding in terms of skill-mastery. Again, the interest of skills lies in the explanation of the mental.

In a recent exchange between Dreyfus and McDowell, Dreyfus argues that in our lives there is an autonomous level of non-conceptual “embodied coping” that we share with other animals. Coping skills – of the special sort discussed by Dreyfus – are said to underwrite a certain “access” to facts in the world *without* the mastery of concepts. Here the relation between skills and cognition is not mediated, as it were, by concepts. Dreyfus’ emphasis is on the coping skills’ ability to open up the world to us in its distinctive, non-conceptual way. He does not have a story about the *relation* between the non-conceptual coping-skills and the upper-storey conceptual life, where explicit knowledge claims are made.

The claim made in this essay about skills that is distinctive from all these authors, is that the mastery of certain kind of skills provides epistemic warrant for knowledge claims. This role of skill-mastery can be compared with perception in that both provide *non-inferential epistemic warrant*. The comparison gives rise *perforce* to the question: is the mastery of skills a kind of *experience*?

Skills and Empiricism

One way to argue that we *need* to conceive of the mastery of skills as a kind of experience is to ask what must be done to the empiricism influenced by Kant for it to have the resource for recognizing the empirical reality of laws of nature?

Kant thought famously that all necessities are rooted in the *rules* of conceptual activities. This is a central move in the “Copernican” turn that was supposed to save laws from Hume’s skeptical dialectic. For reasons I have spelled out in various places of this essay³⁰³, I do not think Kant’s strategy can work. This strategy did have some unfortunate historical legacies. Early Wittgenstein, Carnap, as well as Sellars, all adopted variations of this modality-as-rules theory. In the case of nomological modal statements, Carnap and Sellars thought that they express, directly or indirectly, inference rules of a specific sort. This sort of non-factualism positions about nomological statements make it harder, not easier, to counter the Humean skepticism.

³⁰³ C.f. especially §6.2.2.

Epilogue

The reason why Kant's transcendental argument – though it works for first-order empirical predicates – does not work against Hume's skeptical arguments for *law statements*, is roughly this. Unlike a tree or a house, a law of nature is not a particular: it is general. Consequently knowledge claims about laws cannot directly acquire perceptual warrant, even if the *concept* of law is, according to Kant, necessarily involved in our experience for transcendental reasons. If, on the other hand, warrant for nomological knowledge claims is mediated by perceptual warrant for singular knowledge claims, *some* appeal to *inferential warrant* is needed. But Hume's argument shows that appeal to inferential warrant generates a problem of induction.³⁰⁴

What this suggests is that, what we need, within the framework of a Kantian empiricism, is to enlarge our conception of *experience*, so that there can be *direct experiential warrant* for nomological claims. Mastery of skills discussed in this essay provides an excellent candidate for this kind of non-perceptual experience.

Kant's conception of experience, however, is the conception of something like a collaboration between the faculty of receptivity and the faculty of spontaneity. For skill-mastery to count as experience in this Kantian sense, the notion of receptivity and spontaneity must be substantially modified. Kant's notion of receptivity is a capacity to be *affected*. Experimental success and mastery of skill are hardly cases of being "affected". But there *is* a sense in which they have a receptive component: through experimental success (*nomological*) *information is taken up* ("received"). On the side of spontaneity, Kant's conception of activity is in terms of *conceptual activity*. This is also unduly restrictive. Skills to bring about one thing by doing another is "active" primarily in a non-conceptual, motoric sense. Concepts however do play a role in skills, in the verbal formulation of means and ends, for example.

To accommodate the idea of skill-as-experience in a Kantian empiricism, therefore, though not looking impossible, requires quite a bit of complicated conceptual clarifications. But given the fruits it promises to bear, the work just might be worth it.

³⁰⁴ C.f. §6.2.2 for more details of this argument.

Bibliography

- Anscombe, G. : (1971) "Causality and Determination", Inaugural lecture at Cambridge University, Cambridge University Press, Cambridge. Reprinted in *Metaphysics and the Philosophy of Mind*, pp. 133-148, Basil Blackwell Publisher, 1981.
- Anscombe, G. : (1957) *Intention* (2 ed.). Oxford: Basil Blackwell.
- Au, T. K. : (1984) "Counterfactuals: In reply to Alfred Bloom", in *Cognition* **17**, pp. 289-302
- Barwise, J., & Etchemendy, J.: (1989) "Model-theoretic semantics". In M. Posner (Ed.)
- Bird, A. : (2002) "Laws and Criteria" in *Canadian Journal of Philosophy* **32** pp.511–41.
- Block, N. : (1998) "Conceptual Role Semantics." In *The Routledge Encyclopedia of Philosophy*, edited by Edward Craig. Routledge.
- Bloom, A. H. : (1981) *The Linguistic Shaping of Thought: A Study in the Impact of Language on Thinking in China and the West*. Hillsdale. Erlbaum Associates.
- Brandom, R. : (1998) *Making it Explicit*, Harvard University Press.
- Bretano, F. : (1955) *Psychologie vom empirischen Standpunkt*, ed. Oskar Kraus, 2 vols. Hamburg: Felix Meiner.
- Carroll, J. : (1991) "Property-level Causation?" *Philosophical Studies* **63**: 245-270.
- Cartwright, N. : (1979) "Causal laws and effective strategies", *Noûs*, **13** (4): 419–437.
- _____ : (1980) 'Do the Laws of Physics state the Facts', *Pacific Philosophical Quarterly*, **61**, 75-84.
- _____ : (1999) *The Dappled World*, Cambridge University Press, Cambridge, 1999.
- Church, A. : (1996) *Introduction to Mathematical Logic*. Princeton, NJ: Princeton University Press.
- Collingwood, R.: (1940) *An Essay on Metaphysics*. Oxford: Clarendon Press.
- Davidson, D. : "Causal Relations", *Journal of Philosophy* Volume **LXIV** No. 21. Reprinted in *Essays on Actions and Events*, pp. 149-162. Oxford: Clarendon Press, 1980.
- Davis, W. : "Probabilistic Theories of Causation", in Fetzer (1988), pp. 133-160.
- Dennett, D. : (1989) *The Intentional Stance*, MIT Press, reprint edition.
- Divers, J. : (2006) "Possible-Worlds Semantics Without Possible Worlds: The Agnostic Approach." *Mind* **115** (April): 458.
- Dummett, M. : (1973) *Frege's Philosophy of Language*, New York, Harper and Row.
- _____ : (1993) *The Logical Basis of Metaphysics*, Harvard University Press.
- Eells, E. : (1991) *Probabilistic Causality*. Cambridge, U.K.: Cambridge University Press.
- Evans, Gareth : (1973) "The Causal Theory of Names." *Aristotelian Society Supplementary Volume* **xlvi** : pp.197-200.

- Feng, G. & Yi, L. : (2006) "What If Chinese Had Linguistic Markers for Counterfactual Conditionals? Language and Thought Revisited", in *Proceedings of the Twenty-Eighth Annual Meeting of the Cognitive Science Society*, pp.1281-1286.
- Fetzer, J. (ed.): (1988) *Probability and Causality: Essays in Honor of Wesley C. Salmon*. Dordrecht: Reidel.
- Fodor, J. : (1998) *Concepts: Where Cognitive Science Went Wrong*, Oxford University Press.
- Fodor, J., & Lepore, E.: (2001) "Brandom's Burdens: Compositionality and Inferentialism." *Philosophy and Phenomenological Research* **LXIII**, no. 2 (September).
- Forbes, G. : (1989) *Languages of Modality*. Vol. 9. Basil Blackwell.
- Frege, G. : (1892) "Über Sinn und Bedeutung." *Zeitschrift für Philosophie und philosophische Kritik* **NF 100**: 25-50.
- Gasking, D. : (1955) "Causation and Recipes", *Mind*, **64**, pp. 479–487.
- Geach, P. T. : (1965) "Assertion" *The Philosophical Review* **74**, no. 4: 449-465.
- Gibbard, A., Harper, W.L.: (1978) "Counterfactuals and Two Kinds of Expected Utility", in *IFS*, Harper, Pearce, Stalnaker, eds.
- Goodman, N. : (1947) "The Problem of Counterfactual Conditionals", *Journal of Philosophy* **44**, 113-28.
- Hacking, I. : (1982) "Experimentation and Scientific Realism", *Philosophical Studies*, **61**.
- Harper, W.L., Pearce, G.A., Stalnaker, R. (eds): (1980) *IFS: Conditionals, Belief, Decision, Chance and Time* (The Western Ontario Series in Philosophy of Science), Springer.
- Haugeland, J. : (2000) *Having Thought*, Harvard University Press.
- Hausman, D. M.: (1986) "Causation and Experimentation", *American Philosophical Quarterly*, Volume **23**, No. 2.
- Higginbotham, J.: (1986) "Linguistic Theory and Davidson's Program in Semantics", in E. LePore (ed.) *The Philosophy of Donald Davidson: Perspectives on Truth and Interpretation* (Oxford: Blackwell's), pp. 29–48
- Horwich, P. : (1998) *Truth*, 2nd Edition, Oxford, Oxford University Press, 1998
- _____ : (1998) *Meaning*, Oxford, Oxford University Press, 1998
- Hume, D. : [1739] (1967) *A Treatise of Human Nature*, L. A. Selby-Bigge (ed.), Oxford: The Clarendon Press.
- _____ : [1748] (1955) *An Enquiry Concerning Human Understanding*. Indianapolis: Bobbs-Merrill.
- Jacquette, D. : (1994) "Tarski's Quantificational Semantics and Meinongian Object Theory Domains." *Pacific Philosophical Quarterly* **75**: pp.88-107.
- Janssen, T.M.V.: (1997) "Compositionality", in J.F.A.K. van Benthem & A. ter Meulen (eds.), *Handbook of Logic and Linguistics*, Elsevier Science Publishers, Amsterdam.
- Jeffery, R. C. : (1965), *The Logic of Decision*, McGraw-Hill, New York

Bibliography

- _____ : (1983) *The Logic of Decision*, 2nd ed., University of Chicago Press
- Kaplan, D. : (1989) "Demonstratives." In Almog, Perry, and Wettstein *Themes from Kaplan*. Oxford: Oxford University Press. 1989, pp. 481–563
- _____ : (2004) The Meaning of 'Ouch' and 'Oops.' Howison Lecture in Philosophy, UC Berkeley, August 2004. Transcript (retrieved on July the 13th 2014): <https://sites.google.com/site/2011personindexicals/documen/Kaplan%2C1997-TheMeaningofOuchandOops.pdf>
- Kaufmann, S. : (2004) "A modal analysis of expressive meaning: German ja under quantifiers", Paper presented at Kōbe Shōin University, Kyoto.
- Kitcher, P. & Salmon, W.: (1987) "Van Fraassen on Explanation", *The Journal of Philosophy* **84**, 315-330.
- Kratzer, A. : (1977) "What 'Must' and 'Can' Must and Can Mean", *Linguistics and Philosophy* **1** (3), pp. 337-355.
- _____ : (1989) 'An investigation of the lumps of thought'. *Linguistics and Philosophy* **87** (1), pp. 3–27.
- _____ : (1999) "Beyond *ouch* and *oops*: How descriptive and expressive meaning interact. A comment on David Kaplan's paper", Cornell Conference on Theories of Context Dependency.
- Kripke, S. A. : (1963) "Semantical Considerations on Modal Logic", in Linsky 1971, pp. 63–87. Originally published in *Acta Philosophica Fennica*, **16**.
- _____ : (1972) *Naming and Necessity*. Cambridge, Mass.: Harvard University Press, 1980 (Reprints 1972)
- _____ : (1979) "A puzzle about belief", in A. Margalit (ed.), *Meaning and Use*, Reidel, Dordrecht (239-283)
- _____ : (1982) *Wittgenstein on Rules and Private Language*, Cambridge, Mass.: Harvard University Press.
- Kuhn, T. : (1970) *The Structure of Scientific Revolutions*. 2nd Edition. Chicago: University of Chicago Press.
- Ladymann, J. : (2004) "Constructive Empiricism and Modal Metaphysics: A Reply to Monton and Van Fraassen", *British Journal for the Philosophy of Science*, **55**, pp. 755-765.
- Lewis, David : (1973) *Counterfactuals*, Oxford: Basil Blackwell.
- _____ : (1980) "Index, Context, and Content", in Stig Kanger & Sven Öhman [ed.] *Philosophy and Grammar* (1980; Reidel).
- _____ : (1979) "Attitudes De Dicto and De Se", *The Philosophical Review*, Vol. **88**, No. 4: pp513-543
- McDowell, J. : (1982) "Criteria, Defeasibility, and Knowledge", *Proceedings of the British Academy* **lxviii**, 455-79
- _____ : (1997) "Brandom on Representation and Inference", *Philosophy and Phenomenological Research*, Vol. **LVII**, No. 1
- _____ : (1998) "Having the World in View: Sellars, Kant, and Intentionality" (the Woodbridge Lectures for 1997), *Journal of Philosophy* **xcv**, 431-91

Bibliography

- _____ : (1998) "In Defence of Modesty", Reprinted in *Meaning, Knowledge, and Reality*, pp. 7-107, Harvard University Press.
- _____ : (1998) "Another Plea for Modesty", Reprinted in *Meaning, Knowledge, and Reality*, pg. 108-131, Harvard University Press.
- _____ : (1998) "Reply to Commentators", *Philosophy and Phenomenological Research*, Vol. **LVIII**, No. 2 June 1998.
- _____ : (2005) "Motivating Inferentialism: Comments on Making It Explicit (Ch. 2)", *Pragmatics and Cognition* **13** (1), pp. 121-140.
- _____ : (2001) *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.
- Mellor, D. H. : (1988) "On Raising the Chances of Effects," in *Probability and Causality: Essays in Honor of Wesley C. Salmon*, edited by James Fetzer (Dordrecht: Kluwer), pp. 227-239.
- _____ : (1998) "Necessities and Universals in Natural Laws" in *Philosophy of Science: the Central Issues*, edited by J.A. Cover and Martin Curd, W. W. Norton & Company.
- Menzel, C. : (1990) 'Actualism, Ontological Commitment, and Possible World Semantics', *Synthese* **85**: 355-389.
- Menzies, P. & Price, H.: (1993) "Causation as a Secondary Quality", *British Journal for the Philosophy of Science*, **44**, 187-203, Printed in Great Britain.
- Montague, R. : (1970) "English as a Formal Language". in *Linguaggi nella società e nella tecnica*, edited by B. Visentini, pp.189–223.
- _____ : (1970) "Universal grammar", *Theoria*, 36: 373–398. Reprinted in Thomason (ed.) pp.7-27, 1974
- Nozick, R. : (1969) "Newcomb's Problem and Two Principles of Choice", in *Essays in Honor of Carl G. Hempel*, edited by Nicholas Rescher (Dordrecht: Reidel), pp.114-146.
- O'Shea, J. : (2007) *Wilfrid Sellars: Naturalism with a Normative Turn*. Polity Press.
- Papineau, D. : (1985) "Causal Asymmetry", *British Journal for the Philosophy of Science*, **36**, pp. 273-89.
- _____ : (1991) "Correlations and Causes", *British Journal for the Philosophy of Science*, **42**, pp. 397-412.
- Pelletier, F.J. : (1993) "On an argument against semantic compositionality", in D. Westerståhl, ed., *Logic, Methodology and Philosophy of Science*, Kluwer, Dordrecht.
- Plantinga, A. : (1974) *The Nature of Necessity*. Oxford: Clarendon.
- Potts, C. : (2003) "Expressive content as conventional implicature". In Kadowaki and Kawahara (2003), *Proceedings of the 33rd North East Linguistics Society Meeting*. Amherst, MA: GLSA pp. 303–322
- _____ : (2007) "Expressive content and the semantics of contexts" NSF Grant No. BCS-0642752, <http://www.stanford.edu/~cgpotts/expressives/docs/potts-nsf-expressives-description.pdf>, retrieved on 26 Jan, 2014

Bibliography

- Price, H. : (1991) "Agency and Probabilistic Causality", *British Journal for the Philosophy of Science*, **42**, pp. 157-176.
- _____ : (1986) "Against Causal Decision Theory", *Synthese* **67**, pp. 196-212.
- Priest, G. : (1995) "Etchemendy and Logical Consequence." *Canadian Journal of Philosophy* **25**: 283-292.
- Prior, A. N. : (1962) "The Formalities of Omniscience", in *Philosophy*, Volume **37** Issue 140 Pages 114-129
- Putnam, H. : (1975) "The Meaning of Meaning." *Philosophical Papers, Vol. II : Mind, Language, and Reality*, Cambridge: Cambridge University Press
- Quine, W.V. : (1950) *Methods of Logic*, Holt, Rinehart, and Winston, New York.
- Ramsey, F. P. : (1929) "General Propositions and Causality", in [Ramsey, Frank P., *Philosophical Papers*, Mellor, D.H. (ed.), Cambridge University Press, 1990, 145-163]
- Rooth, M. : (1996) "Focus", In *The handbook of contemporary semantic theory*, ed. Shalom Lappin, 271–297. Oxford: Blackwell.
- Ryle, G. : (1949) *The Concept of Mind*, London: Hutchinson
- Schurz, G. : (1998) "Tarski and Carnap on Logical Truth – or: What is Genuine Logic?" In *Alfred Tarski and the Vienna Circle*, J. Wolenski and E. Köhler, eds., Kluwer.
- Sellars, W. : (1963) 'Empiricism and the Philosophy of Mind', in his *Science, Perception, and Reality*, 1963 London: Routledge and Kegan Paul.
- _____ : (1965) "Scientific Realism or Irenic Instrumentalism", in *Boston Studies in the Philosophy of Science*, edited by Robert S. Cohen and Marx W. Wartofsky, New York.
- _____ : (1969) "Language as Thought and as Communication," *Philosophy and Phenomenological Research* **29**: 506-27.
- _____ : (1978) "The Role of the Imagination in Kant's Theory of Experience." In *Categories: A Colloquium*, edited by Jr. Henry W. Johnstone. Pennsylvania State University Press.
- Soames, S. : (2002) *Beyond Rigidity – The unfinished Semantic Agenda of Naming and Necessity*, Oxford University Press.
- _____ : (2007) "Ambitious Two-Dimensionalism", in *On Sense and Direct Reference*, Matthew Davidson, ed. McGraw Hill.
- Suppes, P. : (1970) *A Probabilistic Theory of Causality*, Amsterdam, North Holland.
- Thomason, R. H. (ed.): (1974) *Formal Philosophy. Selected Papers of Richard Montague*, New Haven: Yale University Press.
- Van Fraassen : (1980) *The Scientific Image*, Oxford University Press.
- _____ : (1990) *Laws and Symmetry*, Oxford University Press.
- von Wright, G.: (1973) *Explanation and Understanding*, Ithaca, New York: Cornell University Press.

Bibliography

- Woodward, J. : (2003) *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
- _____ : (2008) Causation and Manipulability, Entry in Stanford Encyclopedia of Philosophy, last revised on Monday, October the 20th, 2008.
- Wright, C. : (1980) *Wittgenstein on the Foundations of Mathematics*, Cambridge, Mass.: Harvard University Press.