Ilya Dergunov | Christoph Meinerding | Christian Schlag

Extreme Inflation and Time-Varying Expected Consumption Growth

SAFE Working Paper No. 334 | January 2022

Leibniz Institute for Financial Research SAFE
Sustainable Architecture for Finance in Europe
info@safe-frankfurt.de | www.safe-frankfurt.de

Extreme Inflation
and Time-Varying Expected Consumption Growth

Ilya Dergunov[*] Christoph Meinerding[**] Christian Schlag[***]

First version: July 13, 2016
This version: January 3, 2022

Abstract

In a parsimonious regime switching model, we find strong evidence that expected consumption growth varies over time. Adding inflation as a second variable, we uncover two states in which expected consumption growth is low, one with high and one with negative expected inflation. Embedded in a general equilibrium asset pricing model with learning, these dynamics replicate the observed time variation in stock return volatilities and stock-bond return correlations. They also provide an alternative derivation for a measure of time-varying disaster risk suggested by Wachter (2013), implying that both the disaster and the long-run risk paradigm can be extended towards explaining movements in the stock-bond correlation.

Keywords: Long-run risk, inflation, recursive utility, filtering, disaster risk

JEL: E31, E44, G12

[*]Research School of Finance, Actuarial Studies and Statistics, ANU, Acton ACT 2601, Australia. E-mail: ilya.dergunov@anu.edu.au.

[**]Deutsche Bundesbank, Research Centre, 60431 Frankfurt am Main, Germany. E-mail: christoph.meinerding@bundesbank.de.

[***]Faculty of Economics and Business, Goethe University Frankfurt, and Leibniz Institute for Financial Research SAFE, 60323 Frankfurt am Main, Germany. E-mail: schlag@finance.uni-frankfurt.de.

We would like to thank the seminar participants and colleagues at Goethe University, Deutsche Bundesbank, Wharton, BI Oslo, IWH Halle, SNB and OeNB as well as the participants of the WFA 2018 Coronado, SGF 2018, German Economists Abroad Conference 2017, Paris December Finance Meeting 2017, DGF 2017, ESSFM Gerzensee 2017, MFA 2017 , AFFI 2017, and the 2017 Colloquium on Financial Markets in Cologne for their comments and suggestions. Special thanks go to Angela Abbate (discussant), Geert Bekaert, Fernando Duarte, Stefan Eichler (discussant), Philipp Illeditsch, Jan Kragt (discussant), Mattia Landoni, Sydney Ludvigson, Emanuel Mönch, Olesya Grishchenko (discussant), Philippe Mueller, Athanasios Orphanides, Marcel Priebsch, Dongho Song (discussant), Marta Szymanowska, Jules van Binsbergen, Jessica Wachter (discussant), Todd Walker, Rüdiger Weber (discussant), Amir Yaron, and to Franziska Schulte and Iliya Karaivanov for excellent research assistance. A previous version of this paper was circulated under the title "Extreme inflation and time-varying disaster risk". We gratefully acknowledge research and financial support from the Leibniz Center for Financial Research SAFE (formerly Research Center SAFE, funded by the State of Hessen initiative for research LOEWE). This work represents the authors' personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank or its staff.

1 Introduction

Long-run risk asset pricing models rest on the key assumption that the conditional distribution of consumption growth, in particular its mean, is varying over time. Following the seminal publication of Bansal and Yaron (2004), many papers have dealt with the issue of detecting such time variation.[1] In this paper, we take a step back and document the existence of long-run risk by fitting a parsimonious regime-switching model for consumption growth to standard quarterly NIPA aggregate consumption data. We show that the hypothesis of constant expected consumption growth can be rejected at any conventional level of significance.

Though interesting in its own, we document that this fact of time variation in expected consumption growth is more than just a statistical issue. It has implications for the way we think about asset prices: our setup allows us to extend the standard consumption-based long-run risk framework towards explaining the joint dynamics of real and nominal assets. More precisely, augmenting the basic regime switching model by inflation significantly alters the estimated dynamics of the conditional mean of consumption growth. This allows us to match, among other things, the empirically observed time-varying nature of the return correlation between stocks and nominal bonds in a rather simple model featuring recursive preferences and learning. Importantly, we document that a model with time-separable CRRA preferences cannot explain the time-varying (and in particular the episodes of negative) stock-bond correlation.

Our analysis proceeds in several steps. First, we fit a model to US consumption data (starting in 1947) which deviates as little as possible from the assumption of constant expected consumption growth. We find evidence for two regimes in expected growth, and a Wald test strongly rejects the null hypothesis that expected growth is the same in these two states. The importance of this result can hardly be overstated. The asset pricing literature has intensely discussed the pros and cons of long-run risk models. Our test, which is based on

[1] A nice synopsis of arguments against or in favor of long-run risk is given by the two competing papers of Beeler and Campbell (2012) and Bansal, Kiku, and Yaron (2012), respectively. Recent advances were made by Schorfheide, Song, and Yaron (2018), who document and analyze the persistent component in expected consumption growth employing sophisticated Bayesian mixed frequency techniques.

consumption data only, i.e., exclusively on fundamentals, unambiguously supports one key assumption of these models, namely that there is time variation in expected consumption growth.

Second, augmenting this basic model by inflation as a second variable, we detect two states in which expected consumption growth is low, one with very high expected inflation (so-called "stagflation") and one with negative expected inflation ("deflation"). While we abstain from an explicit structural macroeconomic model in our paper, we argue that our dynamics of expected inflation and expected consumption growth are qualitatively in line with the recent New-Keynesian macroeconomic literature on time variation in monetary policy regimes or on the identification of demand and supply shocks. Our model-implied dynamics also match survey data on macroeconomic expectations rather well.

Third, we embed the estimated fundamental dynamics in a standard general equilibrium asset pricing model with recursive preferences and learning. We show that imperfect information about expected consumption growth drives time variation in aggregate stock market volatility and in the stock-bond correlation when we properly condition on inflation as a signal. To this end, we feed the model with the estimated macro dynamics and price stocks and bonds with the model-implied pricing kernel. When comparing model-implied conditional moments with analogous moments from the data, we find a very good fit.

Fourth, we document that both assumptions, recursive preferences as well as learning, are key for this result. In particular, when we feed our estimated fundamental dynamics into a nested model with CRRA preferences, we find that the model has substantial problems to produce a negative stock-bond correlation, and that the model-implied dynamics are strongly at odds with the data in the second half of our sample. The regression of correlations on state variables delivers coefficients with signs opposite to those found in the data. A nested model in which the state of the economy is perfectly observable also has problems to generate the recently observed negative correlation.

Finally, we perform a series of robustness checks and document that our results are robust to constraints on the economic states imposed in the estimation, to subsample analyses, and to variations of the point estimates of the parameters within the error bounds.

What are the economic mechanisms that recursive preferences offer for explaining the stock-bond correlation, and why do CRRA preferences fail here? We basically identify two such mechanisms. First of all, as is well known in the literature, if the representative agent has recursive preferences, the intertemporal substitution effect dominates the income effect upon an exogenous change in expected consumption growth. An increase in real expected consumption growth thus leads to a positive stock return in equilibrium. In a CRRA world, where the income effect dominates the substitution effect (in the usual case of risk aversion exceeding one), the opposite is true, i.e. an increase in real expected consumption growth generates a negative stock return. Second, also well-known, the real risk-free rate is very volatile in standard CRRA settings. Therefore, in a CRRA model realized nominal bond returns reflect both changes in expected inflation and (counterfactually large) changes in real risk-free rates. The magnitudes and the signs of these two channels are much more realistic in our model.

Importantly, our setup is designed to make these different channels visible. We produce all the aforementioned results within a rather simple setup, where no asset price data are used in the estimation of fundamental dynamics. As Schwenkler (2018) and many others have pointed out, this is key for our analysis, since asset price data have the potential to severely confound a macro estimation, e.g., via certain moment conditions to represent the parametric structure of an asset pricing model included in an application of GMM. This is what happens, e.g., in David and Veronesi (2013), who estimate a time series model similar to ours, but do so with moment conditions derived from an asset pricing model with CRRA preferences. The estimation then yields the counterfactual result of expected consumption growth being constant across regimes, because the underlying asset pricing model suffers from the tension between income and substitution effect described above.

The time variation in the correlation between *expected* consumption growth and *expected* inflation generated by our model is also reflective of a broader macrofinance literature. Our model is deliberately silent about monetary policy. Nevertheless, it seems natural to draw a link from our expected inflation and expected consumption growth regimes to persistent changes in monetary policy, for instance during the Volcker disinflation period. Such narratives can be in line with, e.g., New-Keynesian models with time-varying Taylor rule coefficients. Another interpretation could be regime switches in the fiscal/monetary policy mix. Perhaps most closely related to our paper, but more oriented towards monetary economics, is the paper by Song (2017). He distinguishes between procyclical and countercyclical inflation shocks, and between active and passive monetary policy. There is no unique mapping from the states in his model to those in ours, but, e.g., our State 1 (with high expected consumption growth and medium expected inflation) can roughly be identified with a regime that is characterized by countercyclical inflation shocks and passive monetary policy, whereas State 2 (with slightly lower expected growth and substantially lower expected inflation) resembles a regime with rather procyclical shocks and active monetary policy. Similarly, Campbell, Pflueger, and Viceira (2014) attribute changes in the pricing of nominal bonds to a shift towards a more anti-inflationary US monetary policy around 1980 and a shift in the persistence of monetary policy and shocks to the inflation target around 2000. However, these models do not provide a deeper economic understanding of the rather infrequent States 3 and 4 from our estimation. We elaborate more on the macroeconomic background of our model in Section 3.3 below.

In line with papers like Detemple (1986), or Croce, Lettau, and Ludvigson (2012), we make the assumption that the representative investor knows the structural parameters of the model, but cannot observe the true state and thus can only filter the respective probabilities from the data. Given our estimation, extreme (high or low) inflation observations serve as a useful signal, which allows to better infer the time-varying probability of very low or even negative expected consumption growth. The fact that low expected real growth can be linked to high or low expected inflation is key to producing a time-varying stock-bond return correlation. When very high expected inflation occurs together with low expected real growth, both bonds and stocks will tend to have negative returns, resulting in a positive return correlation between the two types of assets. On the other hand, low growth expectations in periods with low expected inflation will lead to negative equity returns, but positive bond returns, implying a negative correlation. Stated differently, the long-run risk paradigm can be extended towards the time-varying nature of the stock-bond return correlation when the signaling role of inflation is taken into account properly.

We test the asset pricing implications of our model by feeding it with the empirical time series of observed consumption growth and inflation and then pricing stocks and bonds using our model-implied pricing kernel. We then compare the dependence of stock return volatilities and stock-bond correlations on the filtered probabilities in the model and the data. In both model and data, the stock-bond return correlation is positively and significantly linked to the probability of being in the high-inflation state, while for the deflationary state the respective regression coefficient is negative, again both in the model and the data. The model also qualitatively matches two key results about stock market volatility. First, volatility is increasing in the (filtered) probability of being in a bad consumption growth state. Second, the relation between volatility and this probability is nonlinear in both model and data.

On the other hand, we have to acknowledge that our model fails to match certain unconditional asset pricing moments in the data. For instance, it does not reproduce the unconditional level of asset price volatility or the dynamic properties of the term structure of interest rates in the data. There are two key reasons for this. First, we strongly tie our hands when estimating the model in the sense that we only rely on fundamental data, but do not make use of any sort of asset price data. As a consequence, the estimation yields a relatively small amount of "long-run risk", i.e., too little persistence in the economic state variables. Second, for the sake of simplicity and tractability, we deliberately shut down channels like stochastic consumption growth volatility which have been shown to be helpful when it comes to solving return volatility puzzles in the literature.

Finally, while our analysis is mainly directed towards time-varying second moments of returns, we find a few interesting results on the side. First of all, our filtered estimates of expected inflation track measures of expected inflation from the Survey of Professional Forecasters (SPF) rather closely.

Second, while there is no comparable measure for expected consumption growth in the SPF, this data set does offer an analogue for the probability of being in a low expected consumption growth state, the so-called "Anxious Index". This index gives the probability that survey participants attach to the event of real GDP growth being negative in the quarter following the quarter in which the survey is taken. We find that our probability estimate derived from consumption and inflation data tracks this index very well.

Third, our model-implied probability for low expected consumption growth also tracks the historical price-earnings ratio of US equity well. It is therefore closely linked to implicit long-run risk variables which have been obtained in the literature through the "reverse engineering" of asset prices . An example of such a variable is the time-varying disaster intensity introduced in Wachter (2013). The correlation between her variable and our filtered probability for being in a low consumption growth states is 0.88, and this co-movement is also reflected in the pronounced correlation between the dividend-price ratios in the data and in our model.

The result concerning the close link between disaster intensity and our filtered probabilities might seem surprising at first glance, since there is no explicit role for disaster risk in our model. However, since the Poisson jumps in the model of Wachter (2013) are uncompensated, changes in their intensity automatically translate into changes in expected consumption growth. Hence, the state variable in her model represents a mixture of both long-run risk and disaster risk. Besides this, papers like Constantinides (2008), Branger, Kraft, and Meinerding (2016), and Schwenkler (2018) argue that disaster risk and long-run risk are intertwined in historical data and hard to disentangle empirically. The episodes of high marginal utility which Wachter (2013) labels as episodes of elevated disaster risk are characterized by low expected consumption growth according to our Markov chain estimation. Our paper thus offers an alternative interpretation of time-varying disaster risk.

We close the paper with a number of robustness checks concerning both the estimation and the asset pricing model. First, we show that a Markov switching model for consumption growth only (i.e., without inflation) cannot replicate the time series of the dividend-price ratio. Second, our results do not depend on expected consumption growth being particularly low (high) in one of the two bad (good) states. A model where expected growth is constrained to be the same in the two good and in the two bad states, respectively, delivers qualitatively the same results. Third, additional analyses with subsamples, with GDP growth instead of consumption growth, or with monthly consumption data also largely confirm our findings.[2]

[2]We also estimate the Markov chain model on a long sample of annual data and identify a deflationary state there as well. As is clear from the discussion above, the identification of a deflationary state is crucial

2 Related Literature

Our paper contributes to and links two major strands of literature, namely the asset pricing literature about inflation as a priced risk factor and the asset pricing literature featuring long-run risk and its empirical estimation. We do not aim at giving a full review of the numerous papers in the latter field. Contributions there were made by Bansal and Yaron (2004), Bansal, Kiku, and Yaron (2016) Beeler and Campbell (2012), Bansal, Kiku, and Yaron (2012), Constantinides and Ghosh (2011), Ortu, Tamoni, and Tebaldi (2013), and Schorfheide, Song, and Yaron (2018), among many others.

Hasseltoft (2009) proposes a model with recursive utility, inflation, and long-run risk to explain (among other things) the time-varying stock-bond return correlation. The mechanism is very different from the one in our paper, since the model in Hasseltoft (2009) heavily relies on stochastic volatility of inflation and consumption growth, but does not allow for a "deflation state", which we clearly identify from the data. In that model, a negative stock-bond return correlation can be generated when inflation volatility decreases and becomes less dominant for risk premiums in general. In such periods, other factors (such as, e.g. dividend growth, which affects only equity prices) gain in relative importance for the comovement of equity and bond prices and may send the stock-bond return correlation into negative territory. However, this mechanism fails to explain the prolonged period of negative stock-bond correlation that we see in the data since roughly 2000, in a period of relatively high inflation volatility.

In a follow-up paper, to address this problem, Burkhardt and Hasseltoft (2012) propose an extension of the model in which the parameters governing the expected consumption growth and expected inflation processes are driven by a two-state Markov chain. The model does better in matching the stock-bond correlation, but, given the way in which the authors introduce inflation risk premia, the asset pricing results seem to a certain degree hardwired into the model. We consider our approach substantially less restrictive in terms of the specification of inflation and consumption growth. Finally, Piazzesi and Schneider (2006) discuss

to match the asset price data. In our benchmark estimation, this identification largely rests on the deflation observations during the Financial Crisis in 2008. However, there are also deflationary episodes in pre-war data, which are not contained in our benchmark sample of quarterly data starting in 1947.

the role of inflation as a signal about future consumption growth, but they focus on the term structure of (nominal and real) interest rates and do not address time variation in the stock-bond correlation.

Song (2017) studies an endowment economy model with recursive preferences, a regime-switching Taylor rule, and a time-varying inflation target. Campbell, Pflueger, and Viceira (2020) analyze the stock-bond correlation in a New Keynesian production economy with habit formation preferences and monetary policy regimes. Complementary to our paper, Constantinides and Ghosh (2017) assess the ability of several macroeconomic predictor variables to improve the performance of consumption-based equilibrium asset pricing models, and they find that inflation data helps to generate the (non-)predictability of price-dividend ratios by cash flow growth rates. In all these papers, however, asset price data is used to calibrate or estimate the model. Ermolov (2018) estimates an external habit model with macroeconomic data only, but his explanation of time variation in the stock-bond correlation is very different from ours. He assumes that consumption and inflation can be hit by two different shocks labeled as supply and demand shocks, whereas we rely on different regimes for expected growth rates.

David and Veronesi (2013) propose Markov switching dynamics for fundamentals, but they assume a model featuring a representative agent with time-additive CRRA preferences who suffers from bounded rationality in the form of money illusion in the spirit of Basak and Yan (2010). Our contribution relative to this paper is twofold. First, David and Veronesi (2013) assume that expected consumption growth is constant, i.e., does not vary with the state of the economy. We show that this assumption is strongly rejected by the data. This fact is important in its own right, but its relevance is not limited to that of a purely statistical result. Besides the dynamics of fundamentals, the second important ingredient to a consumption-based asset pricing model are preferences. As stated above, David and Veronesi (2013) assume CRRA utility (augmented by money illusion), while in our model the representative agent is equipped with recursive preferences. This is not just a matter of taste, but has far-reaching implications for the ability of the model to explain the data. Once one takes the fact into account that expected consumption growth is time-varying, a CRRA model (irrespective of whether it features money illusion or not) cannot explain the pattern in the time series of stock-bond correlations. To sum up, our paper documents that one can take the estimates of the structural parameters purely from fundamentals and still build a powerful asset pricing model, as long as preferences are flexible enough to properly handle the dynamics estimated from fundamentals.

Boons, de Roon, Duarte, and Szymanowska (2020) provide empirical evidence for inflation risk being priced in the cross-section of stock returns. Their paper can be viewed as complementary to ours, since the market price of inflation risk estimated from the cross-section of stock returns switches sign and is linked to the stock-bond correlation in the data. It can be considered a stylized fact that the correlation between inflation and other variables can change the sign of the stock-bond correlation, and we provide a model-theoretic explanation for this result. Other papers in this area include Schmeling and Schrimpf (2011), Balduzzi and Lan (2016), Campbell, Sunderam, and Viceira (2017), Hasseltoft (2012), Ang and Ulrich (2012), and Marfe (2015), to name just a few. Baele, Bekaert, and Inghelbrecht (2010) empirically analyze the determinants of the stock-bond return comovement. Fleckenstein, Longstaff, and Lustig (2016) study the pricing of deflation risk using market prices of inflation-linked derivatives. Applying a wide range of empirical tests, Duffee (2018a,b) argues that a large fraction of the time series variation in nominal yields is explained by variation in real yields, as opposed to variation in expected inflation. He argues that this is at odds with standard asset pricing models like, e.g., those of the long-run risk type. As we show in our analysis, incomplete information and Bayesian learning, by inducing time variation in the correlation between expected inflation and real yields, help to at least partly resolve this issue.

Through the re-interpretation of our state variables as measures of time-varying disaster risk, our paper is also related to this area of research. Rietz (1988) and Barro (2006, 2009) rationalize a high equity premium in the disaster risk framework. Extensions of their basic model are studied by Chen, Joslin, and Tran (2012) and Julliard and Ghosh (2012), among others. Constantinides (2008) criticizes that historically consumption disasters rather unfold over several years instead of just one point in time. The assumption of extreme jumps is also questioned by Backus, Chernov, and Martin (2011). As a response, Branger, Kraft, and Meinerding (2016) combine disaster risk and long-run risk and show that the equity premium puzzle can still be solved with multi-period disasters. Similarly, Gabaix (2012), Wachter (2013), and Tsai and Wachter (2015) analyze models with time-varying jump intensities and recursive preferences. Our results imply that the disaster risk paradigm may be extended towards an explanation of the time-varying stock-bond return correlation, when the effect of inflation on real asset prices is captured properly. Finally, in this regard, our paper may also contribute to the discussion about "dark matter" in asset prices started by Chen, Dou, and Kogan (2020) in the sense that a large fraction of this dark matter may be attributed to uncertainty about extreme inflation.

3 Fundamental Dynamics

3.1 Consumption and inflation

The two fundamental sources of risk in our model are aggregate consumption and inflation. In the baseline version without inflation, we assume that log aggregate real consumption, $\ln C$, follows the process

$$d \ln C_t = \mu^C(S_t)dt + \sigma^C dW_t^C. \tag{1}$$

W^C is a standard Wiener process, the volatility σ^C is constant. The conditional drift rate $\mu^C(S_t)$ is stochastic and follows a continuous-time Markov chain whose current state is denoted by S_t. There are n states (indexed by $i = 1, \dots, n$), with state-dependent drifts μ_i^C. The state transitions of the Markov chain are governed by counting processes whose intensities are collected in the $(n \times n)$-matrix $\Lambda = (\lambda_{ij})_{i,j=1,\dots,n}$. Following the usual convention, we set the diagonal elements λ_{ii} of this matrix equal to $-\sum_{j \neq i} \lambda_{ij}$ so that the rows of Λ sum to zero. In our benchmark empirical case, we will have $n = 2$.

In the full model, the joint dynamics of log aggregate real consumption and the log price level π are given as

$$\begin{aligned} d\ln C_t &= \mu^C(S_t)\,dt + \sigma^C\left(\sqrt{1-\rho^2}\,dW_t^C + \rho\,dW_t^\pi\right) \\ d\pi_t &= \mu^\pi(S_t)\,dt + \sigma^\pi\,dW_t^\pi. \end{aligned} \tag{2}$$

Here W^C and W^π are the (independent) components of a standard bivariate Wiener process. The dynamics in (2) imply that the increments to $\ln C$ and to π are correlated with correlation parameter ρ. The volatilities σ^C and σ^π are assumed constant. The conditional drift rates $\mu^C(S_t)$ and $\mu^\pi(S_t)$ now follow a bivariate continuous-time Markov chain whose current state is again denoted by S_t. Keeping the rest of the notation as above, the number of states in the full model will later turn out to be $n = 4$.

We will often use the vector representation of the above dynamics, which can be written as

$$\begin{pmatrix} d\ln C_t \\ d\pi_t \end{pmatrix} = \mu(S_t)\,dt + \Sigma\,dW_t$$

with

$$\mu(S_t) = \begin{pmatrix} \mu^C(S_t) \\ \mu^\pi(S_t) \end{pmatrix}, \quad \Sigma = \begin{pmatrix} \sigma^C\sqrt{1-\rho^2} & \sigma^C\rho \\ 0 & \sigma^\pi \end{pmatrix}, \quad dW_t = \begin{pmatrix} dW_t^C \\ dW_t^\pi \end{pmatrix}.$$

3.2 Markov chain estimation

To estimate the dynamics of the fundamentals we use quarterly real consumption growth rates from NIPA and quarterly inflation rates constructed according to the Piazzesi and Schneider (2006) mechanism.[3] Our sample period ranges from 1947Q1 to 2014Q1 and represents the longest period for which quarterly data are available.[4] Figure 1 shows time series

[3] The choice of this inflation time series is in line with the literature on consumption-based asset pricing with a focus on inflation risk, with papers like Song (2017), David and Veronesi (2013), and Burkhardt and Hasseltoft (2012) as examples. For a detailed discussion of this issue, we refer the reader to Piazzesi and Schneider (2006).

[4] We have performed the estimation also with alternative samples to compare our findings to those stated in other papers. These results are discussed in Section 6.1. Besides, we have also estimated various constraint versions of the models in which the number of parameters is reduced. This does not change any of our results qualitatively. Details on these constraint models are presented in Section 6.2.

plots of the data.

Based on these data for consumption and inflation we estimate the two models (1) and (2) using maximum likelihood.[5] We assume a constant variance-covariance matrix and only allow for time-varying drifts. Instead of the matrix of transition intensities Λ, the estimation gives us an $(n \times n)$-matrix $Q = (q_{ij})_{i,j=1\dots,n}$ of transition probabilities, which are linked to the intensities via $\lambda_{ij} = -\log(1 - q_{ij})$ for $j \neq i$. The diagonal elements q_{ii} of the transition probability matrix are set such that the rows sum to 1. Standard errors for the parameter estimates are computed via a standard block bootstrap with a block length of ten quarters[6] with overlapping blocks and 5,000 repetitions.

The results for the univariate case are presented in Table 1. The first important finding is that, based on the Bayes Information Criterion (BIC), the algorithm clearly identifies two regimes with values for expected consumption growth of 0.585 and -0.45 percentage points per quarter, respectively. Moreover, a Wald test based on the bootstrapped standard errors strongly rejects the null hypothesis that these two values are equal with a p-value of 0.003. We regard this test as clear evidence that the conditional mean of consumption growth is time-varying, which is the first key contribution of our paper.

Following the seminal publication by Bansal and Yaron (2004), many papers have dealt with the issue of detecting such time variation and verifying this key assumption of the long-run risk model. Most prominently, Beeler and Campbell (2012) argue that the long-run risk model fails to match a series of asset pricing moments in the data, for instance it overstates the predictability of stock returns for consumption and dividend growth and for stock return volatility. Bansal, Kiku, and Yaron (2012) reply to this criticism and provide a set of arguments in support of long-run risk models, for instance by fitting various ARMA models to combined pre- and post-war US consumption data. In a similar spirit, Schorfheide, Song, and Yaron (2018) document and analyze the persistent component in expected consumption growth employing sophisticated Bayesian mixed frequency techniques.

We contribute to this literature by reducing the debate to its very core. Our Markov chain specification with two states for expected consumption growth should be regarded as

[5]For details about the estimation procedure, see Online Appendix A.

[6]Varying the block length does not affect the results.

the smallest possible deviation from a model with constant expected growth, and the statistical test rejects the latter at any conventional significance level. We have further repeated the test with both subsamples of the quarterly post-war data and longer samples of annual data, with our key finding being confirmed in all cases. We regard this as independent, stand-alone evidence that the key assumption of the long-run risk model, namely the existence of non-negligible time variation in expected consumption growth, is strongly supported by the data and can be regarded as a stylized empirical fact.

Adding inflation to the model, the estimation results change significantly (see Table 2). For the bivariate model we identify four regimes: high expected growth–medium expected inflation (state 1), medium expected growth–low expected inflation (state 2), low expected growth–high expected inflation (state 3), and negative expected growth–negative expected inflation (state 4). The estimated transition probabilities imply that state 1 lasts for around 11 quarters on average, while the average time spent in state 2 is 33 quarters.[7] The other two states are not very persistent with an average occupation time of around 6 and 3 quarters, respectively. Most of the time, the economy is thus in state 1 or 2, but it is the rare states 3 and 4 which are very important in the context of asset pricing, since they feature low (or even negative) expected consumption growth. State 3 with high expected inflation and low expected real growth can be labeled "stagflation", whereas in state 4 the expected change in the price level is negative, i.e., there is deflation on average. With a p-value of 0.005, the Wald test again strongly rejects the null hypothesis of equal expected consumption growth across states.

Table 3 reports unconditional moments of consumption and inflation in the data and in the estimated time series model represented by system (2) above. Note that the maximum likelihood estimation does not explicitly target these moments. Still, the general fit of the time series model is good. Only the autocorrelations of consumption growth in the model are a little too low, due to the rather simple Markov chain structure.

A final plausibility check for our results is provided in Figures 2 and 3, where we plot the time series of estimated expected consumption growth and expected inflation resulting from

[7]Transforming probabilities into intensities via $\lambda_{ij} = -\log(1 - q_{ij})$, it follows that the average time spent in state i is given by $-\frac{1}{\sum_{j=1,j\neq i}^{n} \log(1-q_{ij})}$.

our Markov chain estimation, next to the analogous measures from the Survey of Professional Forecasters (SPF). As one can see from these graphs, our expectations time series fit the SPF time series pretty well, in particular for inflation.

3.3 Macroeconomic interpretation

We want to briefly discuss possible macroeconomic interpretations of our time series model. In particular, the time variation in the correlation between *expected* consumption growth and *expected* inflation, which our model produces, deserves further attention. In traditional New-Keynesian macroeconomic models, business cycle fluctuations are largely driven by demand shocks. Such shocks push prices and output in the same direction, and we should expect mostly positive comovement between *realized* output growth and *realized* inflation. Equating output with consumption, State 1 of our estimated model could then ex-post be interpreted as the result of a sequence of positive demand shocks, whereas State 2 would reflect a rather calm period without such shocks. A deflationary episode like State 4 could, in principle, emerge from a sequence of strongly negative demand shocks. Indeed, during the recent period with interest rates at or even below the zero lower bound, policymakers were worried that households and firms expect such a streak of negative demand shocks, leading into a so-called deflationary spiral. A high average inflation state like State 3 in our model could only be justified by a sequence of very large supply shocks. In a recent paper, Ermolov (2018) incorporates this reasoning of demand versus supply shocks and the resulting time-varying comovement of *realized* consumption growth and inflation in an asset pricing model with habit formation preferences.

Expected consumption growth has been studied intensively in long-run risk models for (real and nominal) asset pricing. In a number of such papers, e.g., in Piazzesi and Schneider (2006) and Wachter (2006), the authors argue that a negative correlation between realized inflation and expected consumption growth is generally necessary to generate an upward-sloping nominal term structure. Such a negative correlation may arise in our model when State 3 becomes more likely. According to this line of argument, investors require a premium for holding nominal bonds if positive inflation surprises (i.e., negative bond returns) are associated with low expected consumption growth.

Turning the focus towards *expected* inflation, clearly the stance of monetary policy comes into play. Our model is deliberately silent about monetary policy, but it seems natural to draw a link from our expected inflation regimes to (highly) persistent changes in monetary policy. For instance, Erceg and Levin (2003) develop a DSGE model in which households disentangle persistent and transitory shifts in the monetary policy rule through filtering. Their calibrated model matches the dynamics of output and inflation, in particular through the Volcker disinflation period. Cogley and Sbordone (2008) make a similar point, albeit without learning. Their paper includes an estimate of trend inflation that resembles our expected inflation in Figure 2. Rudebusch and Swanson (2012) analyze the asset pricing properties of a DSGE model with recursive preferences and a nominal long-run risk component, in the sense that the central bank's long-run inflation objective may vary over time. They show that their model can reproduce many stylized facts about equity and term premia.

Adding fiscal policy to the picture, Bianchi and Ilut (2017) analyze regime switching in the fiscal/monetary policy mix. They argue that the fiscal authority was in the lead in the 1960s and 1970s, with only an accommodating role for monetary policy, whereas the relation turned around in the 1980s. According to the authors, it is indeed a regime switch that explains the events in the early 1980s, as opposed to a sequence of shocks. The fiscal/monetary narrative is also in line with earlier studies on New-Keynesian models with time-varying Taylor rule coefficients like Clarida, Gali, and Gertler (2000) and Lubik and Schorfheide (2004). The paper of Duffee (2018a) can also be seen in this tradition. He analyzes to what extent shocks to expected inflation can explain shocks to nominal bond yields. Time variation in the fiscal/monetary policy mix is also at the heart of the "stepping on a rake" narrative that has been prominently discussed by Sims (2011). In an environment of uncertainty about future fiscal policy, policy-generated increases in the nominal interest rate increase rather than reduce inflation.

This type of idea is extended in the direction of asset pricing by Song (2017), who distinguishes between procyclical and countercyclical inflation shocks and between active and passive monetary policy. There is no unique mapping from the states in his model to those in ours, but, for instance, State 1 in our model can roughly be identified with a regime characterized by countercyclical inflation shocks and passive monetary policy, whereas State 2 resembles a regime with procyclical shocks and active monetary policy. Similarly, Campbell, Pflueger, and Viceira (2014) attribute changes in the pricing of nominal bonds to a shift towards a more anti-inflationary US monetary policy around 1980 and and a shift in the persistence of monetary policy and shocks to the inflation target around 2000. These models, however, do not have much to say regarding the rather rare States 3 and 4 from our estimation.

4 Asset Pricing Model

Long-run risk models with time-varying expected consumption growth were designed to explain the time series behavior of asset prices and returns. In the following, we will therefore embed the dynamics estimated in the previous section into a state-of-the-art asset pricing model with recursive preferences and learning. Having inflation in the model allows us to analyze the returns of stocks and nominal bonds jointly. The quantitative results presented in Section 5 are based on the parameter estimates from the benchmark specification in Table 2.[8]

4.1 Preferences

The economy is populated by an infinitely-lived representative investor with stochastic differential utility as introduced by Duffie and Epstein (1992b). The investor has the indirect utility function

$$J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}) = E_t \left[\int_t^\infty f(C_s, J(C_s, \widehat{p}_{1s}, \ldots, \widehat{p}_{ns})) ds \right], \tag{3}$$

where the aggregator f is given by

$$f(C, J) = \frac{\beta C^{1-\frac{1}{\psi}}}{\left(1 - \frac{1}{\psi}\right) \left[(1-\gamma) J\right]^{\frac{1}{\theta}-1}} - \beta \theta J.$$

[8]Alternative parametrizations are discussed in Section 6 below.

In (3), $\widehat{p}_{jt}$ denotes the probability of being in state j which the agent has to filter from the data (see next section). γ, ψ, and β represent the degree of relative risk aversion, the elasticity of intertemporal substitution (EIS), and the subjective time preference rate, respectively, where we we assume $\gamma \neq 1$ and $\psi \neq 1$. Furthermore, $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$. The special case of time-separable CRRA preferences is represented by the restriction $\gamma = \psi^{-1}$, which implies $\theta = 1$. Throughout the paper, we will set $\gamma = 10$, $\psi = 1.7$, and $\beta = 0.02$.[9] With this parameter choice, the agent has a preference for early resolution of uncertainty, since $\gamma > \psi^{-1}$.

4.2 Filtering

We assume that the representative agent cannot observe the current state S_t (and thus does not have perfect knowledge about $\mu^C(S_t)$ and $\mu^\pi(S_t)$). Instead she has to filter her estimates from the data.[10] Incomplete information generates an additional layer of uncertainty that will turn out to be highly relevant for our results. Besides the risk to *switch* to a bad state next period, which would also be present in a full information model, we add the uncertainty about the current regime and thus about the *probability* of switching to a bad regime.[11]

Mathematically, there are two filtrations, $\mathcal{F}$ and $\mathcal{G}$, where $\mathcal{F}$ is generated by the processes $(C_t)_t$, $(\pi_t)_t$ and $(S_t)_t$, whereas $\mathcal{G} \subset \mathcal{F}$ is generated by the processes $(C_t)_t$ and $(\pi_t)_t$ only. The conditional expectations of the drifts given the investor's information, $\widehat{\mu}_t^C$ and $\widehat{\mu}_t^\pi$, are given as

$$\widehat{\mu}_t^C = E\left[\mu^C(S_t)|\mathcal{G}_t\right] = \sum_{i=1}^n \widehat{p}_{it}\mu_i^C$$

and

$$\widehat{\mu}_t^\pi = E\left[\mu^\pi(S_t)|\mathcal{G}_t\right] = \sum_{i=1}^n \widehat{p}_{it}\mu_i^\pi.$$

Here $\widehat{p}_{it} = E\left[\mathbf{1}_{\{S_t=i\}}|\mathcal{G}_t\right]$ denotes the subjective conditional probability of being in state i at time t. As indicated by the structure of the indirect utility in (3), these conditional

[9] A nested version of our model with CRRA preferences ($\theta = 1$) is discussed in Section 5.5 below.

[10] As pointed out in the introduction and as it is standard in the literature, we assume that the representative agent knows the structural parameters of the model. The only thing she does not know is the current state of the economy.

[11] In Section 5.6 we compare our results to those from a nested model with full information, in which the agent can observe the economic state at any point in time.

probabilities represent the state variables in our economy. Since probabilities always sum up to 1, we will have $n-1$ state variables $\widehat{p}_1, \ldots, \widehat{p}_{n-1}$, whose support is the standard simplex in $\mathbb{R}^{n-1}$.

Consumption growth and inflation realizations are observable and serve as a signal for the aggregate state. The dynamics of $\widehat{p}_{it}$ follow from the so-called Wonham filter and are given by

$$d\widehat{p}_{it} = \left(\lambda_{ii}\widehat{p}_{it} + \sum_{j\neq i} \lambda_{ji}\widehat{p}_{jt}\right) dt + \widehat{p}_{it} \left[\begin{pmatrix} \mu_i^C \\ \mu_i^\pi \end{pmatrix} - \sum_{j=1}^{n} \widehat{p}_{jt} \begin{pmatrix} \mu_j^C \\ \mu_j^\pi \end{pmatrix} \right]' (\Sigma')^{-1} \begin{pmatrix} d\widehat{W}_t^C \\ d\widehat{W}_t^\pi \end{pmatrix} \tag{4}$$

with the "subjective" Brownian motions

$$\begin{pmatrix} d\widehat{W}_t^C \\ d\widehat{W}_t^\pi \end{pmatrix} = \Sigma^{-1} \left[\begin{pmatrix} \mu_i^C \\ \mu_i^\pi \end{pmatrix} - \sum_{j=1}^{n} \widehat{p}_{jt} \begin{pmatrix} \mu_j^C \\ \mu_j^\pi \end{pmatrix} \right] dt + \begin{pmatrix} dW_t^C \\ dW_t^\pi \end{pmatrix}.$$

A proof of the filtering equation based on Theorem 9.1 of Liptser and Shiryaev (2001) and a discussion of its properties are provided in Online Appendix B.

In the context of our analysis it is essential to note that the update in the estimated probability $\widehat{p}_i$ depends on both signals, i.e., on both realized consumption growth and realized inflation. Inflation observations have an impact on the perceived probability of being in state i and thus on the conditional expected consumption growth rate. This will be the key driver for our asset pricing results described below.

Figure 4 shows the filtered estimates for the probabilities of the four states, i.e., the estimates the investor would have computed based on information up to and including time t. These estimates are the key quantities analyzed in the following subsection. They will also serve as the explanatory variables in our regression analyses in Section 5. First of all, there is considerable variation in each of the four time series, i.e., the probability of being in state i changes substantially over time. State 1 with the highest expected real growth rate, but also above-average inflation is considered most likely by the investor during the 1960s and much of the 1970s. The investor furthermore perceives a high probability to be in the regime 2 with low inflation and stable growth for extended periods during the 1950s and much of the 1990s, but this probability is very low during the 1970s. Not surprisingly, there is a very high probability for the high inflation state 3 during the latter period. The deflation state 4 is seen as very likely in the beginning of the sample right after the war and during the Great Recession.

4.3 Real Pricing Kernel and Wealth-Consumption Ratio

As shown in Duffie and Epstein (1992a), the real pricing kernel depends on the log wealth-consumption ratio v and is given by

$$\xi_t = C_t^{-\gamma} e^{-\beta\theta t + (\theta-1)\left(\int_0^t e^{-v_u} du + v_t\right)}.$$

The wealth-consumption ratio $I \equiv e^v$ depends on the estimated expected consumption growth $\widehat{\mu}^C$, and therefore in particular on the estimated probabilities $\widehat{p}_i$. It solves a nonlinear partial differential equation given in Online Appendix C.1. A proof and details concerning the numerical solution using a Chebyshev polynomial approximation are also presented in Online Appendix C.1.

Given a solution for I, the pricing kernel has dynamics

$$\begin{aligned} \frac{d\xi_t}{\xi_t} &= -\beta\theta dt - (1-\theta) I^{-1}\, dt - \gamma \frac{dC_t}{C_t} + \frac{1}{2}\gamma^2 (\sigma^C)^2 dt - (1-\theta) \sum_{i=1}^{n-1} \frac{I_{\widehat{p}_i}}{I}\, d\widehat{p}_{it} \\ &\quad + \frac{1}{2} \sum_{i=1}^{n-1} \sum_{j=1}^{n-1} (\theta-1) \left[\frac{I_{\widehat{p}_i \widehat{p}_j}}{I} + (\theta-2) \left(\frac{I_{\widehat{p}_i} I_{\widehat{p}_j}}{I^2} \right) \right] \sigma_{\hat{p}_i} \sigma'_{\hat{p}_j} dt - \gamma(\theta-1) \sum_{i=1}^{n-1} \frac{I_{\widehat{p}_i}}{I} \sigma_{c,\hat{p}_i}\, dt \end{aligned}$$

with the dynamics of $d\widehat{p}_{it}$ given in Equation (4). Importantly, shocks to the state variables $\widehat{p}_i$ affect the pricing kernel. Since these shocks are themselves driven by both consumption and inflation observations, realized inflation indirectly enters the pricing kernel through the learning mechanism.

4.4 Pricing the Assets in the Economy

We are mainly interested in two types of assets, equity and nominal bonds. Equity is defined as a claim to real dividends. When defining dividends, one has to be careful not to alter the informational setup of the model. Dividends are observable, and if they provided a non-redundant signal about the state of the economy, this would affect the initial filtering problem. Technically, this requires the two systems of equations

$$\begin{pmatrix} d\widehat{W}_t^C \\ d\widehat{W}_t^\pi \end{pmatrix} = \Sigma^{-1}\left[\begin{pmatrix} \mu_i^C \\ \mu_i^\pi \end{pmatrix} - \sum_{j=1}^{n} \widehat{p}_{jt} \begin{pmatrix} \mu_j^C \\ \mu_j^\pi \end{pmatrix}\right] dt + \begin{pmatrix} dW_t^C \\ dW_t^\pi \end{pmatrix}$$

and

$$\begin{pmatrix} d\widehat{W}_t^C \\ d\widehat{W}_t^\pi \\ d\widehat{W}_t^D \end{pmatrix} = \Sigma_*^{-1}\left[\begin{pmatrix} \mu_i^C \\ \mu_i^\pi \\ \mu_i^D \end{pmatrix} - \sum_{j=1}^{n} \widehat{p}_{jt} \begin{pmatrix} \mu_j^C \\ \mu_j^\pi \\ \mu_j^D \end{pmatrix}\right] dt + \begin{pmatrix} dW_t^C \\ dW_t^\pi \\ dW_t^D \end{pmatrix}$$

to yield the same solution for $d\widehat{W}_t^C$ and $d\widehat{W}_t^\pi$. Here the superscript D denotes terms related to dividend dynamics and $\Sigma_*\Sigma_*'$ is the covariance matrix of innovations to $\ln C$, π and $\ln D$.

The above condition for the redundance of dividends is satisfied by assuming

$$d\ln D_t = \bar{\mu}\, dt + \phi\left(\sum_{i=1}^{n}(\mu_i^C - \bar{\mu})\widehat{p}_{it}\right) dt + \phi\sigma^C\left(\sqrt{1-\rho^2}\, d\widehat{W}_t^C + \rho\, d\widehat{W}_t^\pi\right).$$

Similar to Bansal and Yaron (2004), the deviation of the drift from its long-term average $\bar{\mu}$ is levered by a factor of ϕ, and we assume $\phi = 5$.[12]

Let ω denote the log price-dividend ratio. Starting from the Euler equation for the price of the dividend claim, we can apply the Feynman-Kac formula to $g(\xi, D, \omega) \equiv \xi D e^\omega$. This yields

$$\frac{\mathcal{A}g(\xi, D, \omega)}{g(\xi, D, \omega)} + e^{-\omega} = 0,$$

where $\mathcal{A}$ denotes the infinitesimal generator. Using Ito's Lemma, we can translate this equa-

[12]The results concerning the stock-bond correlation remain qualitatively the same when ϕ is set to a lower value.

tion into a PDE for ω as a function of the filtered probabilities $\widehat{p}$. This PDE, together with details regarding its derivation, is given in Online Appendix C.3. We solve this PDE again numerically using a Chebyshev approximation.

A nominal bond pays off one unit of money at maturity T, which, in real terms, is equal to $\exp\left(-\int_t^T d\pi_s ds\right) = \exp(\pi_t - \pi_T)$. The price of a nominal bond at time t is thus equal to

$$B_t^{\$,T} = E_t\left[\frac{\xi_T}{\xi_t}\exp(\pi_t - \pi_T)\right].$$

Equivalently, one can define the nominal pricing kernel as $\frac{\xi_T^{\$}}{\xi_t{}^{\$}} \equiv \frac{\xi_T}{\xi_t}\exp(\pi_t - \pi_T)$ and rewrite the pricing formula as

$$B_t^{\$,T} = E_t\left[\frac{\xi_T^{\$}}{\xi_t{}^{\$}}\right].$$

The dynamics of the nominal pricing kernel then follow from Ito's lemma:

$$\frac{d\xi_t^{\$}}{\xi_t^{\$}} = \frac{d\xi_t}{\xi_t} - d\pi_t + \frac{1}{2}\,d[\pi]_t - \frac{d[\xi,\pi]_t}{\xi_t}.$$

Importantly, the nominal risk-free short rate, i.e., the negative of the drift of $\xi^{\$}$, is not just the sum of the real short rate and expected inflation, but involves the covariation between the real pricing kernel and inflation, $d[\xi,\pi]$ (and the quadratic variation of π). The covariation is nonzero if inflation shocks affect the real pricing kernel, as they do in our model, and can be interpreted as an equilibrium inflation risk premium.

The Euler equation and the Feynman-Kac formula applied to $H(\xi_t^{\$}, b_t^{T,\$}) = \xi_t^{\$} e^{b_t^{T,\$}}$ yield a partial differential equation for the (log) price $b_t^{T,\$} \equiv \ln B_t^{T,\$}$ of a nominal zero coupon bond. Details on this partial differential equation and its solution are given in Online Appendix C.5.

5 Results

5.1 Dividend-price ratios

It is instructive to start the analysis by comparing the dividend-price ratio generated by our model with the one observed in the data. To this end, we plug the historical quarterly consumption and inflation time series and our estimated state variables $\widehat{p}_i$ into the numerical solution of the model.

Figure 5 shows the model-implied dividend-price ratio together with the historical dividend-price ratio of the S&P 500 index. The correlation between the two time series is 0.45, and they share all major upward and downward trends. Given that we have not used any asset price data in the estimation, we consider this a pretty strong result. We take this as a first piece of evidence that our proposed state variables $\widehat{p}_i$, embedded into an otherwise standard asset pricing model featuring recursive Epstein-Zin preferences, indeed capture the time variation in valuation ratios well.

5.2 Extreme inflation as a signal about disaster risk

Motivated by the similarity between model-implied and empirical dividend-price ratios, we now turn towards a deeper analysis of the time series pattern of the state variables in our model. The main result of this section is depicted in Figures 6 and 7. The dashed red line in both graphs represents the sum of the filtered probabilities $\widehat{p}_3 + \widehat{p}_4$ from our model. To obtain these filtered probabilities, we plug realized consumption growth and inflation data into our filtering equations and compute the probabilities, which a Bayesian learner would have obtained at each point in time.

We then compare these probabilities to two other time series. The first one is the implied disaster intensity shown in Figure 8 (p. 1017) in the paper of Wachter (2013).[13] She reverse-engineers this quantity from asset prices based on her model, where the intensity of rare consumption disasters follows a mean-reverting process and serves as a state variable. Given

[13] We thank Jessica Wachter for sharing her data with us.

the parameters of her model, she recovers monthly implied values for this state variable from a smoothed time series of historical S&P 500 price-earnings ratios. To make her series comparable to our estimates we take averages of the monthly implied disaster intensities over each quarter. Moreover, the plot shows 5-year moving averages of $\widehat{p}_3 + \widehat{p}_4$ in order to account for the smoothing in Wachter (2013).[14]

The second series is the so-called "Anxious Index" obtained from the SPF. This index, which is available only from 1968 onwards, gives the probability that survey participants attach to the event of real GDP growth being negative in the quarter following the quarter in which the survey is taken. To facilitate the comparison, we lag the Anxious Index by one quarter in the figure.

The two series in Figure 6 exhibit a correlation of 0.88 over our sample period covering almost 70 years, and the correlation in Figure 7 is 0.52. These high levels of correlation are even more remarkable, given that the series are computed from very different data and using rather different approaches. Furthermore, they share all important trends, peaks, and troughs. The time series of Wachter (2013) shows up a pronounced downturn during the 1950s, followed by rather low values in the 1960s, a sharp increase during the 1970s up to around 1982, and then, basically following the same kind of cycle, we observe the sharp decline and low level during the Great Moderation, followed by the recent spike at the beginning of the Great Recession. The SPF time series is more or less reflecting the business cycle, but our probabilities $\widehat{p}_3 + \widehat{p}_4$ also pick up most of the recession periods, except the one in the early 1990s. In particular, both the high inflation regime and the deflation regime and their respective probabilities are relevant here. For instance, a look at the time series plots of the filtered probabilities in Figure 4 shows that the peak of $\widehat{p}_3 + \widehat{p}_4$ in the early 1980's can be traced back to the high probability of the high inflation regime (state 3) over that period, and the deflation regime prevails towards both the beginning and the end of the sample period. This result strongly supports the notion that inflation can serve as a signal for expected real consumption growth in that it allows to quantify the probability of large negative future consumption shocks.

[14]More precisely, for her reverse engineering exercise, she uses the ratio of prices to the previous 10 years of earnings. The two time series depicted in Figure 6 both exhibit an autocorrelation of around 0.99.

To check whether it is indeed inflation that is important here, and not just certain special characteristics of the consumption time series, we redo the analysis based on only consumption data using the univariate baseline model presented in the beginning. Figure 8 presents the time series of the estimated probability for the state with low expected growth. Already from a first rough inspection it becomes clear that the disaster intensity is matched much less precisely than before. The correlation between the series based on Wachter (2013) and the filtered probability of being in the low expected growth state derived from consumption only goes down to roughly 0.33. Even more importantly, the series of estimated probabilities for low consumption growth is now substantially off during basically all periods when the risk of the economy being in a bad state is actually high, e.g., during most of the 1970's and 1980's and also to a certain degree towards the end of the sample period. We interpret these findings as a clear indication that information about inflation in necessary to obtain reliable estimates for the probability of low real growth.

We chose these two time series for a comparison with our model, because their interpretation as "probabilities" (of very low consumption growth) is more in line with our state variables $\widehat{p}_i$ than, for instance, time-varying conditional means of consumption growth, which are typically discussed in long-run risk papers. Nevertheless, besides a simple "sanity check" for our filtered probabilities, the pictures also provide additional insights on the asset pricing literature featuring rare disasters and long-run risk.

Our model does not explicitly feature "disaster risk". Instead, our time-varying expected consumption growth is more in line with a long-run risk interpretation. However, Constantinides (2008), among others, argues that it is empirically very challenging to distinguish models with long-run risk from those featuring time-varying rare disaster risk. Branger, Kraft, and Meinerding (2016) support this argument based on the international panel of consumption growth data used by Barro (2006) and others to calibrate disaster risk models. The authors document that the null hypothesis of these data having been generated by a Markov switching model similar to ours cannot be rejected at conventional significance levels.

Similarly, one can argue that the asset pricing model of Wachter (2013) is not a pure disaster model either, but rather a mixture of long-run risk and disaster model. The expected consumption growth rate in her model is time-varying and a linear function of the stochastic disaster intensity. Consequently, the labeling of the time series plotted in Figure 6 as "implied disaster probability" is disputable. Our comparison shows that it can equally well be interpreted as a long-run risk time series.

5.3 Conditional stock return volatilities

Given that Wachter (2013) documents the important role of time-varying disaster intensities for the dynamics of second moments of returns, we continue the discussion of our asset pricing results by analyzing second moments.

We proceed in the following way. We take the time series of filtered probabilities as shown in Figure 1, plug them into our model solution and compute model-implied real prices for equity and for nominal bonds with five years to maturity. From these time series of real prices we compute model-implied quarterly real log returns for these two assets. More precisely, with S_t and $B_t(20)$ denoting the price of the equity claim and the 20-quarter (five-year) nominal zero coupon bond in quarter t, the returns from quarter t to quarter $t+1$ are computed as $\ln(S_{t+1} + D_{t+1}) - \ln S_t$ and $\ln B_{t+1}(19) - \ln B_t(20)$. We then add log realized inflation to the real returns to obtain nominal returns. The corresponding quantities in the data are quarterly returns of the CRSP value-weighted index and log bond returns computed from the US Treasury yield curve data provided by Gürkaynak, Sack, and Wright (2007)[15] from 1962 on. As the final input to our analyses we compute 20-quarter rolling window return volatilities and correlations and regress them on (the logarithm of) 20-quarter moving averages of the relevant state probabilities $\widehat{p}$. Note that these right-hand variables are the same for model and data in all the regressions reported below.

For the regressions in the model and in the data we state Newey-West adjusted t-statistics with 20 lags, but in addition we also provide confidence intervals derived from a Monte Carlo simulation of the model (shown in square brackets below the respective coefficient). Here we first simulate the model given the dynamics for the fundamentals and

[15] The data are available for download at `http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html`.

the filtered probabilities in Equations (1) to (4) with monthly time increments over a time span of 68 years, corresponding to the length of our sample period for the macroeconomic variables. These monthly data are then aggregated to quarterly and used in the regressions in the same way as described before, i.e., we only use the later 50 years of each sample path, corresponding to the period over which financial market data are available. We repeat this exercise 5,000 times and then take the 5% and the 95% quantile to to obtain 90% confidence intervals.[16]

A look at Figure 9 shows that our model nicely reproduces the patterns of state-dependent stock return volatilities in the data along two important dimensions. First, the estimated probabilities for the states with low consumption growth, $\widehat{p}_3+\widehat{p}_4$, exhibit a positive covariation with stock return volatilities. Second, this relationship is nonlinear and concave, both in the model and in the data. It is worth noting that the second result confirms another prediction from the model of Wachter (2013), namely that stock market volatility is a concave function of time-varying disaster risk.[17] Figure 10 shows the time series of volatilities in the data and in the model.

Motivated by the visual evidence, we regress stock return volatilities on the logarithm of the moving averages of the relevant state probabilities $\widehat{p}$. Table 4 reports the results. The regression coefficients are positive and significant in both model and data. The R^2 is also high in both model and data, and the R^2 obtained from the data is within the simulated confidence bounds for the model.[18] Moreover, almost all of the regression coefficients from the data are within the simulated confidence bounds for the model. The only exception with respect to this last point is the low constant in our model regressions, which indicates that the model-implied unconditional stock return volatility is somewhat on the low side.[19]

Overall, we conclude that the relation between the conditional probability of low con-

[16]Due to the discretization error in the simulation it sometimes happens that the sum of the filtered probabilities exceeds 1 by a very small amount. In that case we rescale the filtered probabilities such that they exactly sum to 1.

[17]See, for instance, Figure 4 (p. 1002) in Wachter (2013).

[18]Not surprisingly, the model R^2 is generally higher than that found in the data, which is due to the stylized structure of the model, where the relations between certain variables are naturally "cleaner" than in the data.

[19]We discuss this issue in detail in Section 6.3 below.

sumption growth (given by the sum $\widehat{p}_3 + \widehat{p}_4$) and stock market volatility is indeed nonlinear, and our model reproduces this stylized fact. Finally, given that our estimation is based on two macro time series only, the goodness of fit in Figure 10 is also remarkable, with a correlation between the data and the model-implied time series of 0.47.

To get an intuition for how our model generates these results, a look at the filtering equation (4) is instructive. The nonlinear expressions in front of the Wiener processes make $\widehat{p}_3$ and $\widehat{p}_4$ fluctuate a lot more when they are at intermediate levels as compared to when they are close to 0 or 1. Moreover, the unconditional averages of both $\widehat{p}_3$ and $\widehat{p}_4$ are rather low, so that higher values for $\widehat{p}_3$ and $\widehat{p}_4$ more or less automatically go together with higher fluctuations in these quantities. States 3 and 4 are the least persistent in our estimation, so that when the agent currently has a high estimate for the probability of being in one of these two states, this probability is likely to decrease again quickly.

Altogether, there is thus more movement in state variables when the likelihood of low consumption growth is large. This in turn means that a high likelihood of low consumption growth goes together with a high volatility of the price-dividend ratio, which is a function of these state variables. In sum, the higher the likelihood of low consumption growth, the higher the equity return volatility. Note that our model of course also reproduces the relevance of the probability of being in a good state, i.e., of the sum $\widehat{p}_1 + \widehat{p}_2$, for stock return volatilities. The coefficients in the data and in the model are both significantly negative, which follows from the simple fact that $\widehat{p}_1 + \widehat{p}_2 = 1 - \widehat{p}_3 - \widehat{p}_4$.

Finally, Table 5 presents the results from additional regressions to investigate the notion of a signaling role of inflation for second moments of stock returns. Here we regress the same left-hand side variable as before on rolling averages of what we call the extreme entropy of the state distribution (defined as $\widehat{p}_3 \ln \widehat{p}_3 + \widehat{p}_4 \ln \widehat{p}_4$). We propose this quantity as another proxy for uncertainty about low consumption growth. We also show results for expected inflation (defined as $\sum_{i=1}^{4} \mu_i^{\pi} \widehat{p}_i$) and realized inflation as explanatory variables.

The results for extreme entropy can be interpreted in the way that uncertainty about extreme inflation and low consumption growth is a main driver of stock return volatilities, and our model provides an economic equilibrium mechanism to explain this stylized fact.

On the other hand, expected and realized inflation as the right-hand side variables do not have explanatory power in the data and are only marginally significant in the model as well (the simulated confidence bounds include zero). This provides additional support for our model. Inflation itself cannot explain stock return volatilities, unless it is decomposed into components capturing the risk of very high inflation (represented by $\widehat{p}_3$) and of a deflationary regime (represented by $\widehat{p}_4$), respectively. Our Markov chain estimation implies that inflation is positively correlated with $\widehat{p}_3$, but negatively correlated with $\widehat{p}_4$. Depending on the amount of observations from the deflation regime along a given sample path, the coefficient from a regression of equity return volatility on inflation in model-generated data can thus be positive or negative.

Recursive preferences are a key ingredient of our model. With respect to this feature, our paper is closely related to recent studies like Benzoni, Collin-Dufresne, and Goldstein (2011) and Drechsler (2013), where it is shown that models featuring recursive preferences, coupled with learning about fundamentals, are very well able to match stylized facts about stock return volatility, in particular in terms of its dynamics (i.e., capturing the predictive power of implied volatilities, variance risk premia, and other related conditional quantities).[20]

5.4 Conditional stock-bond return correlations

For the results concerning stock market volatility, it is the total estimated probability of being in a bad state for expected consumption growth, given by the sum $\widehat{p}_3 + \widehat{p}_4$, which is relevant. When we now look at the stock-bond return correlation, the distinction between the two bad consumption states with respect to expected inflation becomes relevant. Figure 11 shows the time series of correlation in the data and in the model. Tables 6 and 7 contain the results of our regression analyses, Figure 12 presents the corresponding scatter plots.[21]

The most important result here is that both in the model and in the data the estimated coefficient for $\log(\widehat{p}_3)$ is positive (and highly significant), while the coefficient for $\log(\widehat{p}_4)$ is

[20] The impact of preferences is also analyzed in more detail in Section 5.5 below.

[21] Note that we regress 'raw' correlation ρ on the state variables. Since correlation is bounded between -1 and 1, transformations like $\tilde{\rho} = \ln(\frac{1+\rho}{1-\rho})$ might seem warranted to guarantee that the regressions are well specified. Rerunning all our regressions using this transformation leaves the results qualitatively unchanged.

negative (and also highly significant). Given that we do not use any asset price information to estimate the fundamental dynamics in our model, the similarity between model and data appears again remarkable. Moreover, the scatter plots again reveal a pronounced nonlinearity in the relation between correlation and the filtered probabilities, both in the model and in the data. Finally, the goodness of fit in Figure 11 is also very pronounced. The correlation between the data and the model-implied time series is 0.44 in the upper picture.

Our benchmark data is the sample of interpolated term structures of U.S. Treasury bonds provided by Gürkaynak, Sack, and Wright (2007), which is shorter than our macro data sample. Given that the latter starts in 1947, we can in principle compute model-implied asset returns for the whole time period from 1947 to 2014. As a robustness check, we therefore determine the stock-bond correlation in the data from an alternative sample, namely the holding period return of the Ibbotson U.S. Long-Term Government Bond Index which measures the performance of 20-year maturity U.S. Treasury bonds. These index returns are available from 1926 onwards, but since our macro estimation uses quarterly data after 1947, we constrain ourselves to the sample from 1947 until 2014.

The lower graph in Figure 11 shows the resulting time series of rolling window correlations. As one can see from the figure, the model-implied correlation tracks the correlation in the data closely also in this extended sample. The two time series have a correlation of 0.49, i.e., they exhibit the same correlation over this longer sample as over our shorter benchmark sample. Interestingly, in the data we now see another period with negative correlation in the very early part of the sample, and our model captures this negative correlation as well.

The only period where our model-implied stock-bond correlation is slightly off is the early 1990s. As can also be seen from Figure 7, our stylized time series model featuring only consumption and inflation does not capture this recession period in the U.S., suggesting that the economic downturn during this time is driven by factors which are beyond the scope of our model. When we exclude the period 1991-1998 from our sample, the correlation between model-implied and empirical stock-bond correlation goes up even further to 0.66.

What is the mechanism inside the model responsible for these patterns? First, an increase in $\widehat{p}_3$ makes it subjectively more likely for the investor that the economy is in the high inflation state now, and probably also over the next few quarters. This has a positive effect on expected inflation and a negative effect on expected consumption growth across all horizons. In this case the bond return over the next quarter is composed of a positive 'carry' component (which is, if nothing changes, equal to the yield of the bond) and a negative component due to an upward shift in the nominal yield curve. In general the second effect dominates the first, so that bond prices tend to go down. Note that the upward shift in the nominal yield curve itself is the composite of two effects: an increase in expected inflation and a slight decrease in the level of the real yield curve. The second of these two effects is typically negligible with recursive preferences, and therefore, the nominal yield curve shifts upwards in response to an increase in $\widehat{p}_3$. The stock return upon a positive shock to $\widehat{p}_3$ depends on real quantities only. A high $\widehat{p}_3$ implies that the economy is more likely to be in a low consumption growth regime, and stock prices tend to be low in such an environment. Taken together, the reactions of bond and stock prices to an increase in $\widehat{p}_3$ go in the same direction, implying a positive correlation.

State 4 is a low expected inflation state with low expected growth, so the response of bond prices to a high $\widehat{p}_4$ is different. Again, there is the positive carry return. But now there is also an additional positive return because the nominal yield curve shifts downwards in response to a higher probability for deflation. This is because the increase in $\widehat{p}_4$ decreases both expected inflation and expected consumption growth across all horizons. Altogether, the impact of a high $\widehat{p}_4$ on bond returns is large and positive. At the same time, such a high $\widehat{p}_4$ signals a high likelihood of low (even negative) expected consumption growth, which depresses equity prices. In sum, the impact of a high likelihood for the deflationary regime is strong on both stock and bond prices, but it is of opposite signs, implying a negative correlation between the two types of assets.

The above findings concerning the role of $\widehat{p}_3$ and $\widehat{p}_4$ for the conditional stock-bond correlation are well in line with the literature. In a purely empirical paper, Baele, Bekaert, and Inghelbrecht (2010) try to fit the correlations of daily stock and bond returns with a multi-factor model. They find that macro factors (in particular output gap and inflation) do not add much explanatory power when the loadings of stock and bond returns are assumed to be constant over time. The performance of the macro factors improves, however, in a regime switching estimation when the loadings are allowed to switch sign. Our findings are potentially related to this in the sense that we show that the overall risk of low expected consumption growth, proxied by the sum $\widehat{p}_3 + \widehat{p}_4$, does not predict correlation, neither in the model nor in the data (the confidence bounds derived from simulation include zero).

For both model and data, the regressions with extreme entropy generate the expected result with insignificant coefficient estimates. The reason is again that this measure captures general uncertainty about bad consumption growth states. Uncertainty about being in the deflation state, however, decreases correlation, whereas uncertainty about the high inflation state increases it. An aggregate measure of uncertainty cannot capture these two opposing effects adequately. Again, the R^2 values in the model are on the high side due to the somewhat stylized model structure, but the fact that the confidence bounds for the model-based regressions contain the R^2 from the data indicates that our model generates a sensible goodness of fit.

5.5 Recursive Preferences, Money Illusion, CRRA preferences

Recursive preferences are essential to match the time-varying stock-bond correlation in our model. We have also solved a version of our model with constant relative risk aversion preferences, where we set the elasticity of intertemporal substitution to $\psi = \frac{1}{10}$ so that $\theta = 1$, but leave the parameters that we obtain from the Markov chain estimation unchanged.

Figure 13 shows the time series of model-implied return correlations. Remarkably, the model is no longer able to produce a negative stock-bond correlation in the 2000s. The constrained model does not match the historical time series at all, the correlation between model and data is even negative at -0.15. Table 8 presents results from regressions analogous to those presented in Table 6. The constrained model generates positive and significant regression coefficients for both $\widehat{p}_3$ and $\widehat{p}_4$.

To get the intuition behind this result, look again at the three components of the holding period bond return as described in Section 5. Both the carry component and the change in expected inflation are independent of the representative agent's preferences, but the change in the real yield curve is clearly not, since real bond prices are determined in equilibrium. A slightly higher current value of $\widehat{p}_3$ or $\widehat{p}_4$ results in a reduced estimate of conditional expected consumption growth, which leads to a massive decline in the overall level of the real yield curve in a CRRA economy. It is well known that CRRA models have a hard time matching the empirically observed smoothness of the real risk-free rate. Altogether, in a CRRA economy bond returns are thus positively related to both $\widehat{p}_3$, the probability of being in a stagflation regime, and $\widehat{p}_4$, the probability of being in a deflationary regime.

Concerning stock returns, high values for $\widehat{p}_3$ or $\widehat{p}_4$ signal low expected consumption growth, while consumption volatility is not affected. With the usual popular CRRA parameterization of $\gamma > 1$ and, consequently, $\psi = \gamma^{-1} < 1$, the income effect dominates the substitution effect, and a lower expected consumption growth rate actually implies a higher stock price. This means that stock returns will be positively related to both $\widehat{p}_3$ and $\widehat{p}_4$ in a CRRA economy (and the effect is stronger for $\widehat{p}_4$, since State 4 has the lowest expected consumption growth rate). This can also be seen from Figure 14, which plots the model-implied dividend-price ratio in the constrained CRRA model together with the historical dividend-price ratio in the data. The two time series are actually negatively correlated. Altogether, we can conclude that with CRRA preferences the stock-bond return correlation loads positively on both $\widehat{p}_3$ and $\widehat{p}_4$, which is at odds with the data.

The above facts imply that, in order to obtain results similar to ours in a CRRA model, one has to keep the variation in the expected consumption growth rate small enough to mitigate the consequences of the typical counterintuitive CRRA result that prices are low in high-growth states. This is exactly the path taken by David and Veronesi (2013). Since the estimation in David and Veronesi (2013) relies on both asset pricing and macroeconomic data and uses moment conditions from an asset pricing model with CRRA preferences, they find that expected consumption growth is hardly varying across states, so they constrain it to be equal across states in their following numerical evaluation (see p. 703 of their paper). This is in pretty strong contrast to our estimation results presented above. We show that recursive preferences can resolve this contradiction, i.e., with this more flexible utility representation the problem of a counterintuitive link between fundamentals and valuation ratios vanishes.

Moreover, in contrast to David and Veronesi (2013), we also do not need to assume any sort of bounded rationality on the part of the representative investor.[22]

5.6 Full Information

In our model we make the assumption that the current state of the economy is unobservable and has to be filtered from macro data. To see why this imperfect observability is relevant in the context of our analysis, we also solve a version of our model with full information. The theoretical solution is presented in Appendix D, the proof is a slight modification of the proof in Appendix A of Branger, Kraft, and Meinerding (2016). In an economy with four states, we obtain four possible values for the wealth-consumption ratio, for the price-dividend ratio and for the price of a nominal bond with a given maturity.

To evaluate the model with full information, we perform the following exercise, which is analogous to our analyses in the previous sections. Our Markov chain estimation above gives time series of ex-post probabilities that the economy was in one of the four states at a particular historical point in time. We assume that the state with the highest probability is the true state. Combining this time series of states with the model solution gives us historical time series of wealth-consumption ratios, price-dividend ratios, and bond prices under the assumption of full information. From these time series and historical consumption and inflation shocks, we then again compute time series of model-implied returns, and from these model-implied returns we compute rolling-window correlations and volatilities as before.

For brevity, we discuss the stock-bond return correlation results only because the failure of a full information model becomes most evident here. Figure 15 shows the respective time series from model-generated and empirical data. The two time series differ substantially throughout the sample. Most importantly, the model with full information has problems to generate the negative correlation between stock and bond returns that we see in the data in

[22] In addition to the constant expected consumption growth, David and Veronesi (2013) also assume that the representative agent suffers from money illusion to make inflation enter the pricing kernel. The agent (in their setup irrationally) bases real decisions partly on nominal variables. Basak and Yan (2010) show that, with CRRA utility, this assumption results in a pricing kernel which is composed of the original real pricing kernel and an adjustment for inflation. We also solved versions of our model with CRRA preferences and money illusion, but this does not change any of the results discussed in this section.

the most recent years.

The reason for this poor performance can be traced back to the estimated Markov chain transition probabilities (Table 2). According to the estimation, there is a zero probability to enter the deflationary state 4 from state 2 or 3. In an economy with full information, this implies that the threat of entering the worst possible economic state in the next quarter is shut down completely for about 70% of the time. In an economy with partial information, on the other hand, there is always a nonzero subjective probability assigned to being in a state from which the economy can slide into state 4. Consequently, the average risk of being in a state with low expected growth is much lower in the full information economy. Moreover, in the full information economy the price-dividend ratio of the stock is high in states 2 and 3 and low in states 1 and 4 , while it should be high in states 1 and 2 and low in states 3 and 4 in order to match the time variation in the stock-bond correlation.[23]

6 Robustness

6.1 Alternative samples

The finding that extreme inflation provides information about low expected real consumption growth is, strictly speaking, of course based on the estimation for our given sample. To alleviate potential concerns that our results might be sample-specific, we repeat the estimation with four alternative data sets. The first two are subsamples of our quarterly consumption and inflation data starting in 1962 and 1965, respectively, which are used in David and Veronesi (2013) and Burkhardt and Hasseltoft (2012). Moreover, we analyze monthly US consumption and inflation data, which are available from 1959 onwards. Finally, we re-estimate our model with GDP growth rates instead of consumption growth rates, where the former are available on a quarterly basis starting in 1947.

Figure 16 shows the proxies for the time-varying probability of low expected consumption growth obtained by applying exactly the same methodology as before to the alternative

[23]The exact numbers for the price-dividend ratio in the full information model are 138.9, 159.8, 150.5, and 143.2 for states 1 to 4, respectively.

samples. More precisely, we proceed as follows. For every sample, we estimate the time series model as defined in Equations (1) and (2). The number of states identified by the Bayes Information Criterion is four for all samples. As in the benchmark estimation, we label the states in which the conditional expected consumption growth rate μ_i^C is below the unconditional average consumption growth rate as "bad states". This refers to two of the four states in each of the cases. We then compute the filtered probabilities $\widehat{p}_i$ for these two bad states and add them. In the upper pictures in Figure 16, the red dashed line shows this sum. The blue solid line is the same as in Figure 6. Finally, we also repeat the estimation without inflation data, i.e., with only consumption or GDP data. In these cases we have only two states, and we treat the state with the lower expected consumption growth rate as the bad state. The filtered probability of this state is depicted in the lower row of pictures.

We can draw two conclusions from this exercise. First, recovering the time series of implied disaster probabilities from Wachter (2013) is to a very large degree independent of the specific sample. In each of the graphs in Panel A, the red dashed line tracks the blue line very closely, the correlations between the two time series are in fact even higher than for the benchmark sample (0.88, 0.90, 0.89, and 0.89, respectively). Second, the result that the replication fails with consumption (or GDP) data only is also confirmed. The best fit is obtained in the case with GDP instead of consumption data, but the correlation between the two time series is only 0.51. The monthly consumption data is too noisy to replicate the time-varying probability of low expected consumption growth.

6.2 Constrained model specifications

Besides analyzing alternative samples, one might also consider imposing more structure on the Markov chain model. The benchmark specification has 27 free parameters to estimate, so there may be constrained versions of the model in which the number of parameters can be reduced without losing too much explanatory power. One may even allow for more states, but, e.g., restrict expected consumption growth rates or expected inflation to be the same across some of these states. Generally, the number of possible constrained models is infinitely large, and we think that an unconstrained estimation provides the cleanest setup. Nevertheless, given our interpretation of states 3 and 4 as the two bad states, an obvious candidate for a constrained model is one in which expected consumption growth is equal across the two good states and across the two bad states, i.e. $\mu_1^C = \mu_2^C$ and $\mu_3^C = \mu_4^C$. This constraint is also justified by the fact that our estimates of the expected consumption growth rates in the two good states and in the two bad states are relatively close together. For instance, using the bootstrapped sample paths, an F-test yields that the joint hypothesis $\mu_1^C = \mu_2^C$ and $\mu_3^C = \mu_4^C$ cannot be rejected at the 10%-level. Moreover, among the many different constrained models we have estimated, this particular one exhibits the lowest BIC value and should thus be preferred.

Figure 17 presents the proxy for time-varying probability of low expected consumption growth that we obtain in the constrained model. This time series is very similar to those that were generated using the benchmark unconstrained model.[24]

To sum up, the finding that extreme inflation helps to recover the time-varying probability of low expected consumption growth is robust across different samples and robust to the constraint of equal expected real growth rates in the two bad and the two good states, respectively. If anything, the benchmark specification, on which we rely throughout the paper, yields conservative estimates with respect to the correlation between our and Wachter's measure for the probability of being in a bad state for real growth.

6.3 Parameter variations and unconditional moments

The unconditional asset pricing moments generated by our model are computed via the Monte Carlo simulation as described in Section 5. The results are shown in Table 9. When interpreting the numbers, one has to keep in mind that our model is estimated only on the basis of fundamental data for consumption and inflation, i.e., it is not calibrated to match unconditional return moments, and that there are no additional risk factors like stochastic volatility. So it should not come as a surprise that the model does not match the data perfectly with respect to unconditional risk premia or volatilities. The equity premium generated by our model, for instance, is roughly 1 percentage point and the equity return volatility is 7.5

[24]Moreover, none of our asset pricing findings change when we use this constrained specification.

percentage points annually, which reflects the fact that there is relatively little variation in the market prices of risk. The average spread between bonds with a maturity of 5 years and those with 3 months is small on average in the data and in the model (where it is basically equal to zero). Finally, the unconditional stock-bond correlation is matched pretty well by the model.

A few slight modifications of the parametrization lead to substantial improvements with respect to unconditional asset pricing moments. Besides the results for our benchmark parametrization, Table 9 presents unconditional moments from four such variations. Parametrization 2 is the same as the benchmark parametrization except that we lower the expected consumption growth in states 3 and 4 by one standard error, so that now $\mu_3^C = 0.032$ and $\mu_4^C = -0.407$. Such a parametrization reflects results from an estimation with GDP growth instead of consumption growth data. Parametrization 3 differs from the benchmark parametrization in that we increase the probability of the Markov chain to remain in state 3 or 4 by one respective standard error. The new third and fourth row of the transition matrix are then given as $(0.057, 0.027, 0.916, 0.000)$ and $(0.000, 0.137, 0.000, 0.863)$, respectively. This reflects the fact that the bad states (in particular state 4) are relatively rare and their persistence is thus estimated relatively imprecisely (see Table 2). The only change in Parametrization 4 relative to the benchmark is the use of the leverage parameter $\phi = 8$. Finally, the column labeled "Combined" reports the results when Parametrizations 2 and 3 are applied jointly.

In all these modifications, the equity premium is substantially higher than in the benchmark case, reaching about 6 percentage points when Parametrizations 2 and 3 are combined, and also the equity return volatility increases. The largest overall effect on unconditional moments can be seen in Parametrization 3. This is in line with existing research on asset pricing models with long-run risk or disaster risk, in which parameters governing the average length of the time spent in a low expected growth or disaster state have a very large effect on the equity premium (see, e.g., the discussion in Branger, Kraft, and Meinerding (2016)). When calibrating such a model to post-war consumption data only, one may considerably underestimate this length.

On the other hand, even with these modifications, the model falls short in generating a reasonably large asset price volatility. For instance, the real and the nominal risk-free short rate in our model are not volatile and not persistent enough. This is consistent with the general failure of our model to generate a sufficiently high stock price volatility (see also the regression intercepts in Table 4 or the graph in Figure 9). Similarly, the persistence and volatility of the price-dividend ratio (not shown in the table) are rather low.

There are two major reasons for this general failure. First, our estimated Markov chain model exhibits a rather low persistence of consumption growth and inflation (see Table 3). The estimation does not target moments like the autocorrelation explicitly. And since we deliberately choose a very simple time series model, which is estimated with macro data only, expected consumption growth is found to be time-varying, but the persistence is not enough to generate sufficient variation in the tails of the distribution of the equilibrium pricing kernel. Long-run risk models in the literature are typically calibrated/estimated *with the explicit goal* of matching asset pricing moments, i.e. the authors always make a compromise between macro and asset pricing moments, sacrificing a bit of macroeconomic fit to asset pricing purposes.

Second, the conditional volatility of consumption growth and inflation is constant in our model. Similar to Bansal and Shaliastovich (2013), we make this assumption in order to sharpen our understanding of the impact of time variation in expected consumption growth and expected inflation. It is, however, well known from the literature that a pronounced persistence of shocks to macroeconomic volatility strongly helps to match unconditional risk premia and return volatilities. The respective mechanisms have been studied extensively in papers like Bansal and Yaron (2004) or Bansal, Kiku, and Yaron (2016).

The unconditional stock-bond correlation is hovering around zero across all parametrizations shown in Table 9, similar to its value in the data. We would like to emphasize though that a key takeaway from our paper is the existence of pronounced swings in the *conditional* stock-bond correlation, which can take large positive and large negative values both in the model and in the data. This implies that the *unconditional* stock-bond correlation is a relatively uninformative moment when it comes to assessing the quality of the model.

Finally, we also do not do a good job at matching properties of the term structure of interest rates, and in particular of long-term bonds. Again, the reason is that our model is deliberately simple and lacks a lot of features that have been found to be crucial for bond pricing. The macro variables are not persistent enough in our estimated model so that we do not have enough “long-run risk”. This finding is consistent with other papers featuring long-run risk models for bond pricing like Bansal and Shaliastovich (2013). These authors use interest rate data to estimate their model and find that a high persistence of expected consumption growth and expected inflation is key to reproduce the dynamic properties of the term structure. However, as Table 9 shows, there is considerable variation in the slope of the term structure. While the slope is close to zero unconditionally, the conditional term structure in our model can be upward-sloping or downward-sloping, depending on the state variables $\widehat{p}_i$.

7 Conclusion

Long-run risk models for asset pricing rest on the key assumption that the conditional distribution of consumption growth is time-varying. We provide evidence in favor of this assumption using a simple Markov regime switching model for expected consumption growth. While already interesting in itself, this reduced-form approach turns out to be very fruitful when it comes to explaining the joint dynamics of real and nominal asset prices. Augmenting the time series model by inflation as a second macro variable significantly alters the estimated regimes. In particular, we then find two states in which expected consumption growth is low, one with high expected inflation and one with negative expected inflation. Embedding the estimated dynamics in a standard general equilibrium asset pricing model with recursive preferences and learning allows us to match time series of aggregate stock return volatility and the stock-bond return correlation.

The basic intuition underlying our results is that low consumption growth tends to occur together with either very high or very low inflation. In contrast to the volatility of stock returns, where it is mainly the *overall* probability of the two bad states for expected consumption growth that matters, it is the distinction between the two with respect to expected inflation which becomes relevant for the stock-bond return correlation. In the high expected inflation state, stocks and bonds will both tend to have negative returns, so that their correlation will be positive, while in the deflationary state, stocks will still do poorly, but nominal bonds will exhibit positive returns, resulting in a negative correlation between the two types of assets.

Our research design differs from other approaches to calibrate dynamic asset pricing models, where often both asset pricing moments and macroeconomic moments are used to identify the deep parameters of the model. This has the potential to yield a parametrization where macro dynamics are not matched very well. Given the parsimony of the research design, in particular the fact that the model is estimated from two macro time series only, we regard our results as evidence that the long-run risk paradigm in asset pricing can be extended towards an explanation of the time-varying stock-bond return correlation when the signaling role of inflation is properly accounted for. Besides, we also document that our filtered probability of being in a bad consumption growth regime closely tracks the evolution of state variables like the so-called "Anxious Index" from the Survey of Professional Forecasters or the "time-varying disaster risk" which Wachter (2013) obtains via reverse-engineering asset prices. Our findings thus also contribute to the recent debate about the empirical distinguishability of long-run risk and disaster risk models.

In summary, our paper shows that a large part of the variation in asset prices may actually be attributable to inflation risk, and that the long-run risk model class provides a promising framework to study the link between inflation, consumption growth and both real and nominal asset returns.

References

ANG, A., AND M. ULRICH (2012): "Nominal Bonds, Real Bonds and Equity," Working Paper.

BACKUS, D., M. CHERNOV, AND I. MARTIN (2011): "Disasters implied by equity index options," Journal of Finance, 66(6), 1969–2012.

BAELE, L., G. BEKAERT, AND K. INGHELBRECHT (2010): "The Determinants of Stock and Bond Return Comovements," Review of Financial Studies, 23(6), 2374–2428.

BALDUZZI, P., AND C. LAN (2016): "Survey Forecasts and the Time-Varying Second Moments of Stock and Bond Returns," Working Paper.

BANSAL, R., D. KIKU, AND A. YARON (2012): "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," Critical Finance Review, 1(1), 183–221.

BANSAL, R., D. KIKU, AND A. YARON (2016): "Risks for the Long Run: Estimation with Time Aggregation," Journal of Monetary Economics, 82(1), 52–69.

BANSAL, R., AND I. SHALIASTOVICH (2013): "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," Review of Financial Studies, 26(1), 1–33.

BANSAL, R., AND A. YARON (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," Journal of Finance, 59(4), 1481–1509.

BARRO, R. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," Quarterly Journal of Economics, 121(3), 823–866.

——— (2009): "Rare disasters, asset prices, and welfare costs," American Economic Review, 99(1), 243–64.

BASAK, S., AND H. YAN (2010): "Equilibrium Asset Prices and Investor Behavior in the Presence of Money Illusion," Review of Economic Studies, 77(3), 914–936.

Beeler, J., and J. Campbell (2012): "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," Critical Finance Review, 1(1), 142–182.

Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein (2011): "Explaining asset pricing puzzles associated with the 1987 market crash," Journal of Financial Economics, 101(3), 552–573.

Bianchi, F., and C. Ilut (2017): "Monetary/fiscal policy mix and agents' beliefs," Review of economic Dynamics, 26, 113–139.

Boons, M., F. de Roon, F. Duarte, and M. Szymanowska (2020): "Time-Varying Inflation Risk and the Cross-Section of Stock Returns," Journal of Financial Economics, 136, 444–470.

Branger, N., H. Kraft, and C. Meinerding (2016): "The Dynamics of Crises and the Equity Premium," Review of Financial Studies, 29(1), 232–270.

Burkhardt, D., and H. Hasseltoft (2012): "Understanding Asset Correlations," Working Paper.

Campbell, J., A. Sunderam, and L. Viceira (2017): "Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds," Critical Finance Review, 6(2), 263–301.

Campbell, J. Y., C. Pflueger, and L. M. Viceira (2014): "Monetary policy drivers of bond and equity risks," Working Paper.

Campbell, J. Y., C. Pflueger, and L. M. Viceira (2020): "Macroeconomic drivers of bond and equity risks," Journal of Political Economy, 128(8), 3148–3185.

Chen, H., W. Dou, and L. Kogan (2020): "Measuring the "Dark Matter" in Asset Pricing Models," Working Paper.

Chen, H., S. Joslin, and N.-K. Tran (2012): "Rare disasters and risk sharing," Review of Financial Studies, 25, 2189–224.

Clarida, R., J. Gali, and M. Gertler (2000): “Monetary policy rules and macroeconomic stability: evidence and some theory,” The Quarterly journal of economics, 115(1), 147–180.

Cochrane, J. H. (2018): “Stepping on a rake: The fiscal theory of monetary policy,” European Economic Review, 101, 354–375.

Cogley, T., and A. M. Sbordone (2008): “Trend inflation, indexation, and inflation persistence in the New Keynesian Phillips curve,” American Economic Review, 98(5), 2101–26.

Constantinides, G. (2008): “Comment on “Macroeconomic Crises since 1870” by Barro and Ursua,” Brookings Papers on Economic Activity, pp. 341–50.

Constantinides, G., and A. Ghosh (2011): “Asset Pricing Tests with Long-Run Risk in Consumption Growth,” Review of Asset Pricing Studies, 1(1), 96–136.

——— (2017): “What Information Drives Asset Prices?,” Working Paper.

Croce, M., M. Lettau, and S. Ludvigson (2012): “Investor Information, Long-Run Risk and the Term Structure of Equity,” Review of Financial Studies, 28(3), 706–742.

David, A., and P. Veronesi (2013): “What Ties Return Volatilities to Price Valuations and Fundamentals,” Journal of Political Economy, 121(4), 682–746.

Debortoli, D., and R. Nunes (2014): “Monetary regime switches and central bank preferences,” Journal of Money, credit and Banking, 46(8), 1591–1626.

Detemple, J. B. (1986): “Asset pricing in a production economy with incomplete information,” Journal of Finance, 41, 383–391.

Drechsler, I. (2013): “Uncertainty, Time-Varying Fear and Asset Prices,” Journal of Finance, 68(5), 1843–1889.

Duffee, G. (2018a): “Expected Inflation and Other Determinants of Treasury Yields,” Journal of Finance, 73, 2139–2180.

——— (2018b): "Expected Inflation, Real Rates, and Stock-Bond Comovement," Working Paper.

Duffie, D., and L. G. Epstein (1992a): "Asset Pricing with Stochastic Differential Utility," Review of Financial Studies, 5(3), 411–436.

——— (1992b): "Stochastic Differential Utility," Econometrica, 60(2), 353–394.

Ehling, P., M. Gallmeyer, C. Heyerdahl-Larsen, and P. Illeditsch (2018): "Disagreement about Inflation and the Yield Curve," Journal of Financial Economics, 127(3), 459–484.

Eraker, B. (2008): "Affine General Equilibrium Models," Management Science, 54(12), 2068–2080.

Eraker, B., I. Shaliastovich, and W. Wang (2016): "Durable Goods, Inflation Risk and Equilibrium Asset Prices," Review of Financial Studies, 29(1), 193–231.

Erceg, C. J., and A. T. Levin (2003): "Imperfect credibility and inflation persistence," Journal of Monetary Economics, 50(4), 915–944.

Ermolov, A. (2018): "Time-Varying Risks of Nominal Bonds: How Important are Macroeconomic Shocks," Working Paper.

Fleckenstein, M., F. Longstaff, and H. Lustig (2016): "Deflation Risk," forthcoming: Review of Financial Studies.

Gabaix, X. (2012): "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," Quarterly Journal of Economics, 127(2), 645–700.

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007): "The U.S. Treasury yield curve: 1961 to the present," Journal of Monetary Economics, 54, 2291–2304.

Hamilton, J. D. (1994): Time series analysis. Princeton University Press.

Hasseltoft, H. (2009): "The "Fed Model" and the Changing Correlation of Stock and Bond Returns: An Equilibrium Approach," Working Paper.

——— (2012): "Stocks, Bonds, and Long-Run Consumption Risks," Journal of Financial and Quantitative Analysis, 47(2), 309–332.

Ireland, P. N. (2007): "Changes in the Federal Reserve's inflation target: Causes and consequences," Journal of Money, credit and Banking, 39(8), 1851–1882.

Julliard, C., and A. Ghosh (2012): "Can rare events explain the equity premium puzzle?," Review of Financial Studies, 25, 3037–3076.

Kim, C.-j., P. Manopimoke, and C. R. Nelson (2014): "Trend inflation and the nature of structural breaks in the New Keynesian Phillips Curve," Journal of Money, Credit and Banking, 46(2-3), 253–266.

Lansing, K. J. (2009): "Time-varying US inflation dynamics and the New Keynesian Phillips curve," Review of Economic Dynamics, 12(2), 304–326.

Liptser, R., and A. Shiryaev (2001): Statistics of Random Processes I: General Theory. Springer-Verlag, Berlin.

Lubik, T. A., and F. Schorfheide (2004): "Testing for indeterminacy: An application to US monetary policy," American Economic Review, 94(1), 190–217.

Marfe, R. (2015): "Labor Rigidity, Inflation Risk and Bond Returns," Working Paper.

Newey, W. K., and K. D. West (1987): "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," Econometrica, 55(3), 703–708.

Ortu, F., A. Tamoni, and C. Tebaldi (2013): "Long-Run Risk and the Persistence of Consumption Shocks," Review of Financial Studies, 26(11), 2876–2915.

Piazzesi, M., and M. Schneider (2006): "Equilibrium Yield Curves," NBER Macroeconomics Annual, pp. 389–442.

Rietz, T. (1988): "The equity premium: A solution," Journal of Monetary Economics, 22(1), 117–31.

Rudebusch, G. D., and E. T. Swanson (2012): "The bond premium in a DSGE model with long-run real and nominal risks," American Economic Journal: Macroeconomics, 4(1), 105–43.

Schmeling, M., and A. Schrimpf (2011): "Expected Inflation, Expected Stock Returns, and Money Illusion: What can we learn from Survey Expectations?," European Economic Review, 55(5), 702–719.

Schorfheide, F., D. Song, and A. Yaron (2018): "Identifying Long-Run Risks: A Bayesian Mixed-Frequency Approach," Econometrica, 86(2), 617–654.

Schwenkler, G. (2018): "Estimating the Dynamics of Consumption Growth," Working Paper.

Sims, C. A. (2011): "Stepping on a rake: The role of fiscal policy in the inflation of the 1970s," European Economic Review, 55(1), 48–56.

Song, D. (2017): "Bond Market Exposures to Macroeconomic and Monetary Policy Risks," Review of Financial Studies, 30(8), 2761–2817.

Tsai, J., and J. Wachter (2015): "Disaster risk and its implications for asset pricing," Annual Review of Financial Economics, 7, 219–252.

Wachter, J. (2006): "A Consumption-Based Model of the Term Structure of Interest Rates," Journal of Financial Economics, 79(2), 365–399.

Wachter, J. (2013): "Can time-varying risk of rare disasters explain aggregate stock market volatility?," Journal of Finance, 68(3), 987–1035.

Figure 1: Fundamental data

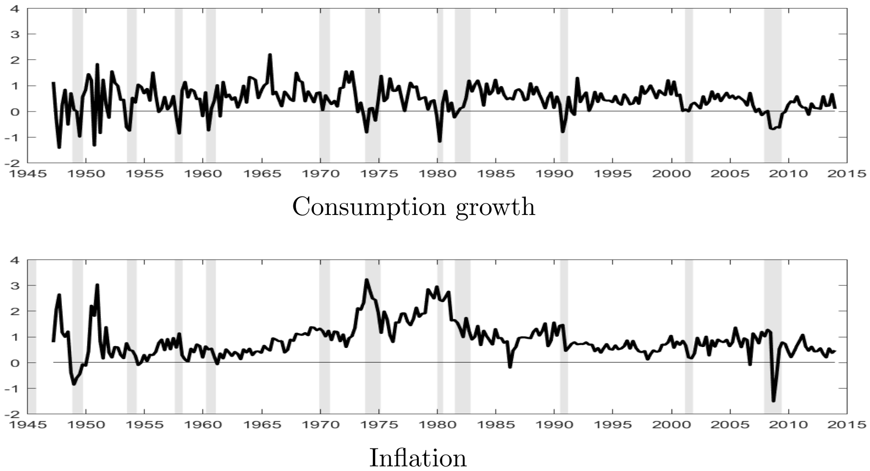

The graph shows time series plots of the quarterly data for consumption growth and inflation over our sample period from 1947 to 2014.

Figure 2: Expected Consumption Growth and Expected Inflation

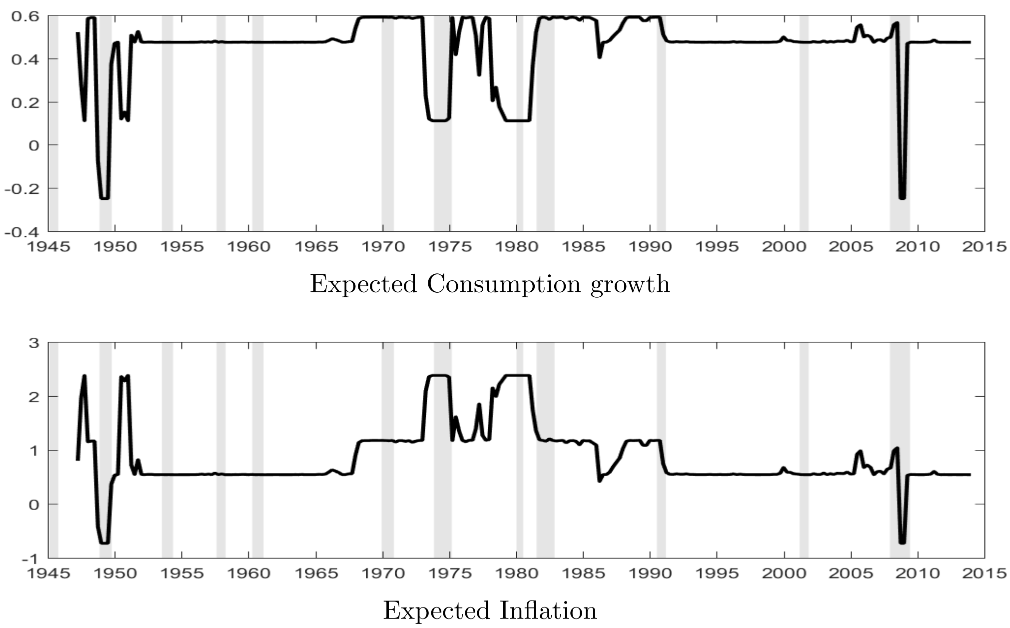

The graphs show time series plots of expected consumption growth and expected inflation obtained from the Kalman filter. Shaded areas indicate NBER recessions.

Figure 3: Expected Inflation and SPF Forecasts

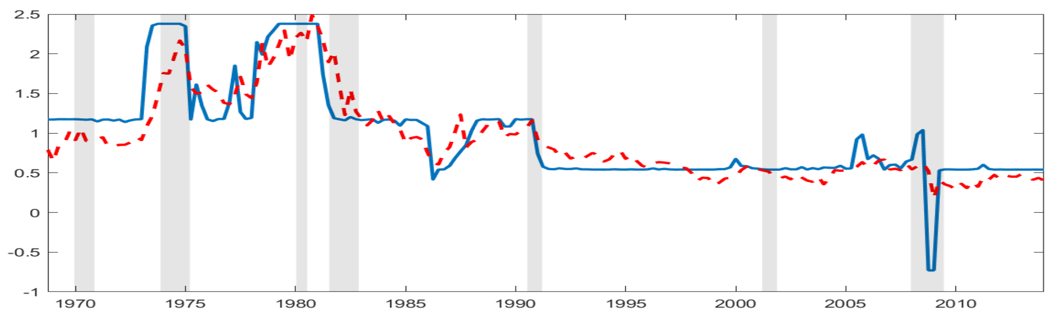

Expected Inflation vs. SPF GDP deflator forecast

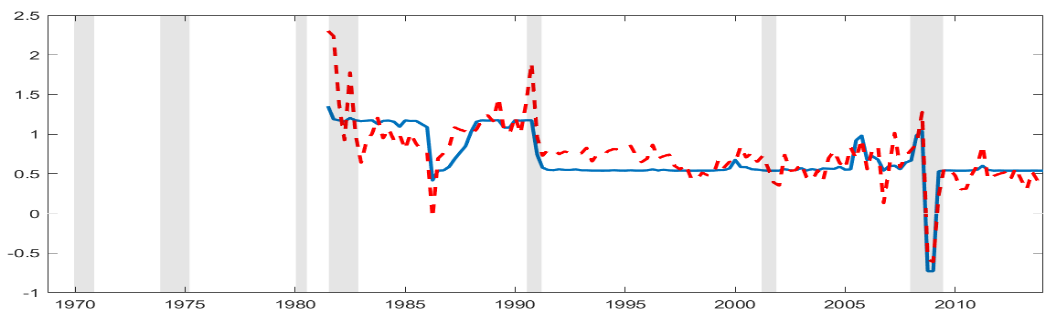

Expected Inflation vs. SPF CPI forecast

The graphs plot expected inflation obtained from the Kalman filter (blue solid line) together with the SPF forecasts for the GDP deflator and the CPI (red dashed lines). Shaded areas indicate NBER recessions.

Figure 4: Real-time filtered probabilities of a Bayesian learner

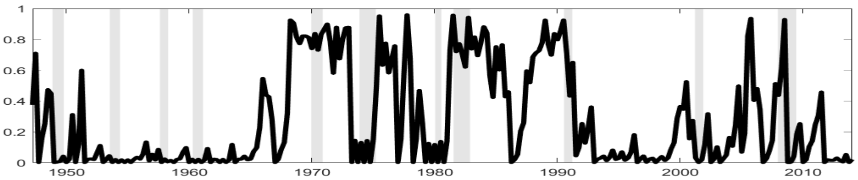

$\widehat{p}_1$

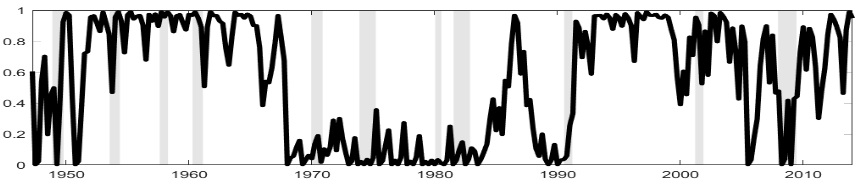

$\widehat{p}_2$

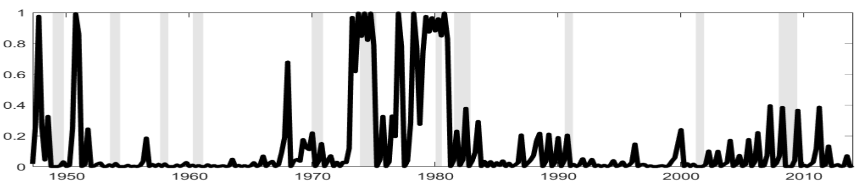

$\widehat{p}_3$

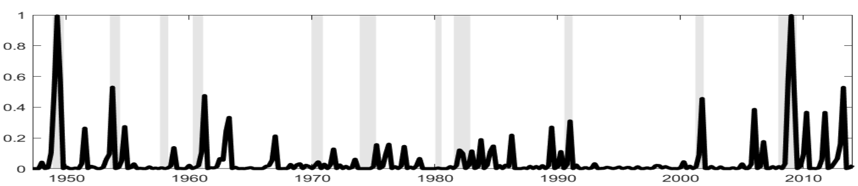

$\widehat{p}_4$

The graphs present the real-time filtered probabilities for each of the four states estimated for our Markov switching models, as defined in Section 4.2, with parameters shown in Table 2. Shaded areas indicate NBER recessions.

Figure 5: Dividend-price ratio

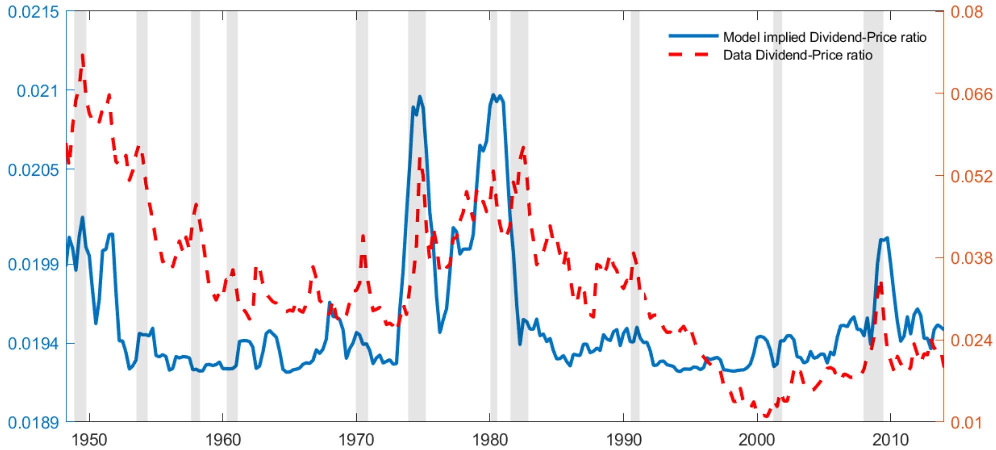

The figure depicts quarterly time series of dividend-price ratios. The blue solid line is the dividend-price ratio of the S&P 500 index. The red dashed line is obtained by plugging the historical paths of consumption, inflation, and our state variables $\widehat{p}_i$ into the numerical solution of the model. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is 0.45.

Figure 6: Time-varying probability of low expected consumption growth vs. implied disaster intensity from Wachter (2013)

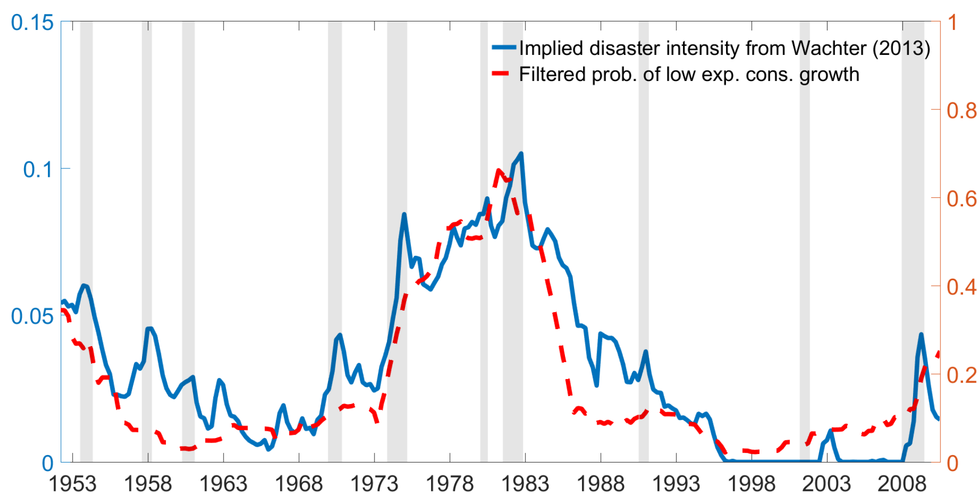

The solid (blue) line depicts the time-varying disaster intensity, which Wachter (2013) extracts from asset price data. The dashed (red) line shows 20-quarter moving averages of the sum of the estimated probabilities $\widehat{p}_3 + \widehat{p}_4$ from our Markov switching model using consumption and inflation data for the period from 1947 to 2014. To obtain $\widehat{p}_3$ and $\widehat{p}_4$ we plug realized consumption growth and inflation data into our filtering equations and compute the probabilities, which a Bayesian learner would have assumed at each point in time. The correlation between the two time series is 0.88.

Figure 7: Time-varying probability of low expected consumption growth vs. Anxious Index

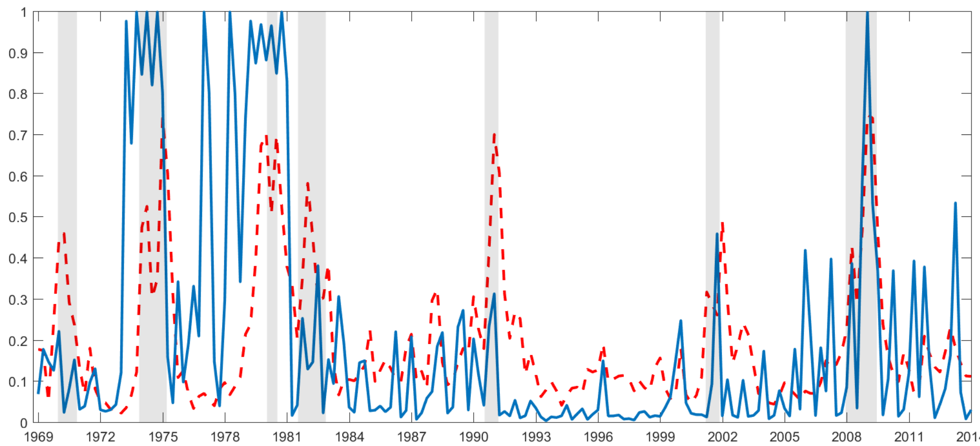

The dashed red line depicts the "Anxious Index", i.e. the probability of a negative realized GDP growth rate in the next quarter, according to the Survey of Professional Forecasters, shifted by one quarter to ease the interpretation of the figure. The solid blue line shows the sum of the estimated probabilities $\widehat{p}_3 + \widehat{p}_4$ from our Markov switching model using consumption and inflation data for the period from 1947 to 2014. To obtain $\widehat{p}_3$ and $\widehat{p}_4$ we plug realized consumption growth and inflation data into our filtering equations and compute the probabilities, which a Bayesian learner would have assumed at each point in time. The correlation between the two time series is 0.52.

Figure 8: Time-varying probability of low expected consumption growth (estimated from consumption growth only)

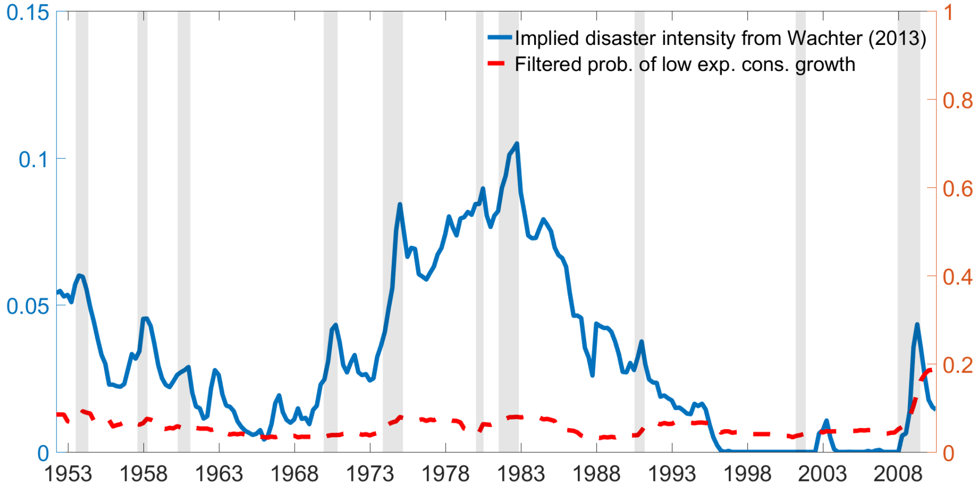

The solid (blue) line depicts the time-varying disaster intensity, which Wachter (2013) extracts from asset price data. The dashed (red) line shows 20-quarter moving averages of the estimated probability for the state with low expected consumption growth from a Markov switching model using only consumption data for the period from 1947 to 2014. To obtain this probability we plug realized consumption growth into our filtering equations and compute the probabilities, which a Bayesian learner would have assumed at each point in time. The correlation between the two time series is 0.33.

Figure 9: Stock return volatilities and probability of low expected consumption growth

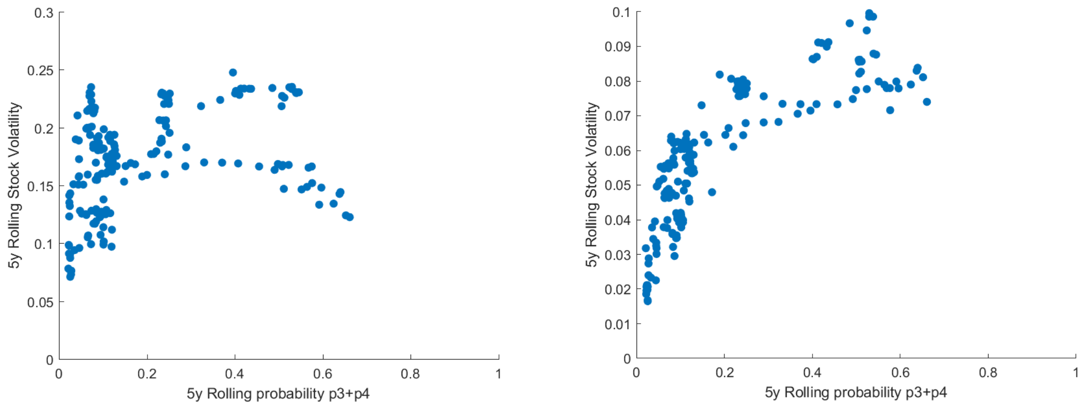

Data Model

The figure depicts scatter plots for the regressions presented in Table 4. The dependent variables in each regression are volatilities of quarterly stock returns computed over rolling windows of 20 quarters. The independent variable is the logarithm of the averages of $\widehat{p}_3 + \widehat{p}_4$ over the same 20 quarters periods. The left figure labeled "Data" is based on the estimated time series of the $\widehat{p}_i$ depicted in Figure 1 as well as the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve (see Gürkaynak, Sack, and Wright (2007)). The right figure labeled "Model" is based on the same time series of the $\widehat{p}_i$, but uses the returns which our model would have implied given this path of consumption, inflation, and the state variables. The financial data for these regressions starts in 1965. The model parameters are estimated using macroeconomic data since 1947.

Figure 10: Conditional stock return volatilities

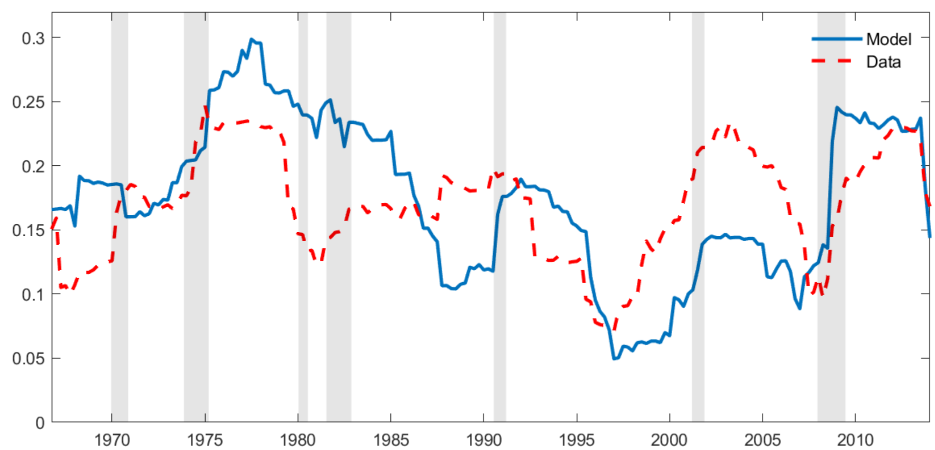

The figure depicts conditional 20-quarter rolling window stock return volatilities. The red dashed line is based on the CRSP value-weighted index. The blue solid line is based on the returns which our model would have implied given the historical paths of consumption, inflation, and state variables (this series is multiplied by 3). The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is 0.47.

Figure 11: Conditional stock-bond return correlation

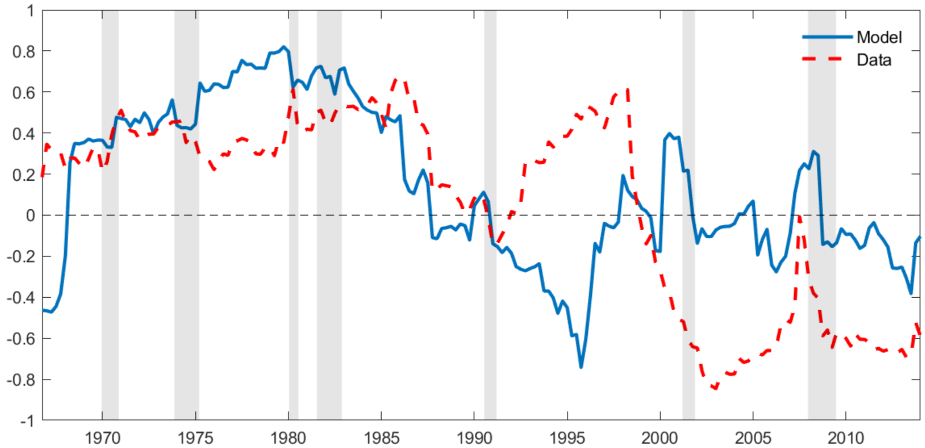

Bond returns based on Gürkaynak, Sack, and Wright (2007) data

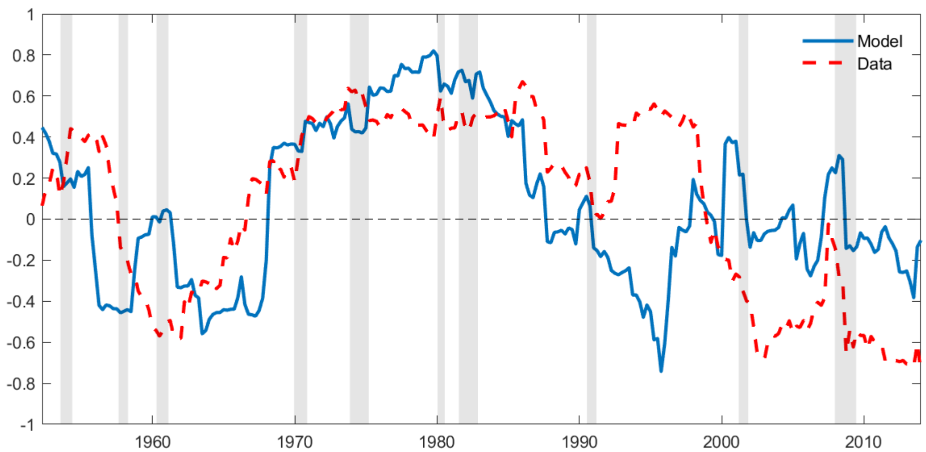

Bond returns based on the Ibbotson U.S. Long-Term Government Bond Index

The figure depicts conditional 20-quarter rolling window correlations of stock returns and returns of 5-year nominal bonds. The red dashed line is based on the CRSP value-weighted index and on the interpolated yield curve data from the Federal Reserve starting in 1962 (upper picture) or the Ibbotson U.S. Long-Term Government Bond Index starting in 1947 (lower picture). The blue solid line is based on the returns which our model would have implied given the historical paths of consumption, inflation, and state variables. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is 0.44 and 0.49, respectively.

Figure 12: Stock-bond return correlations and probabilities for high inflation and deflation

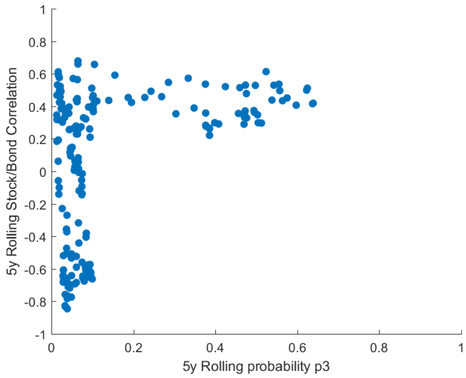

Data
correlation vs. $\widehat{p}_3$

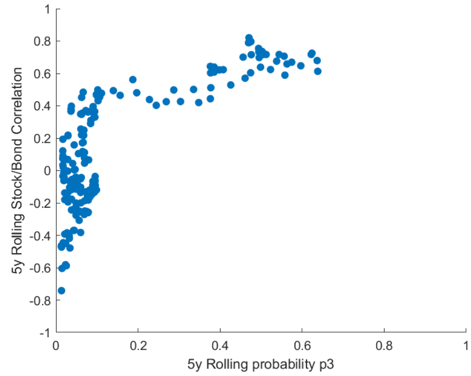

Model
correlation vs. $\widehat{p}_3$

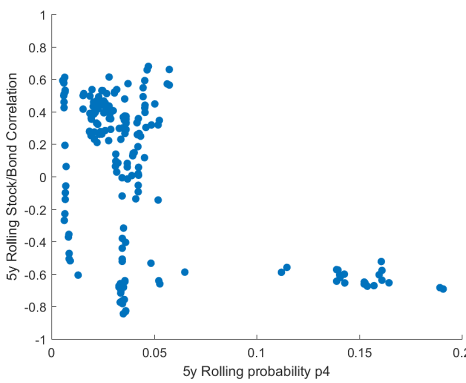

Data
correlation vs. $\widehat{p}_4$

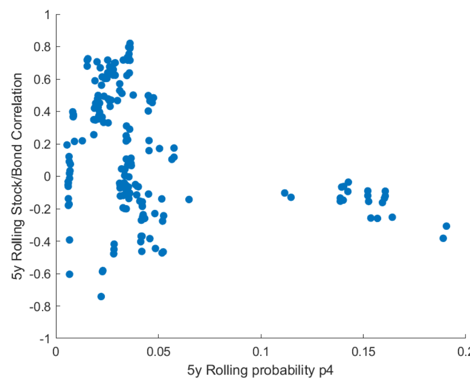

Model
correlation vs. $\widehat{p}_4$

The figure depicts scatter plots for the regressions presented in Table 6. The dependent variables in each regression are correlations of quarterly holding-period returns of stocks and 5-year nominal bonds computed over rolling windows of 20 quarters. The independent variables are the logarithms of the averages of $\widehat{p}_3$ and $\widehat{p}_4$, respectively, over the same 20-quarter periods. The left figures labeled “Data” are based on the estimated time series of the $\widehat{p}_i$ depicted in Figure 1 as well as the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve (see Gürkaynak, Sack, and Wright (2007)). The right figures labeled “Model” are based on the same time series of the $\widehat{p}_i$, but use the returns which our model would have implied given this path of consumption, inflation, and the state variables. The financial data for these regressions starts in 1965. The model parameters are estimated using macroeconomic data since 1947.

Figure 13: Conditional stock-bond return correlation with CRRA preferences

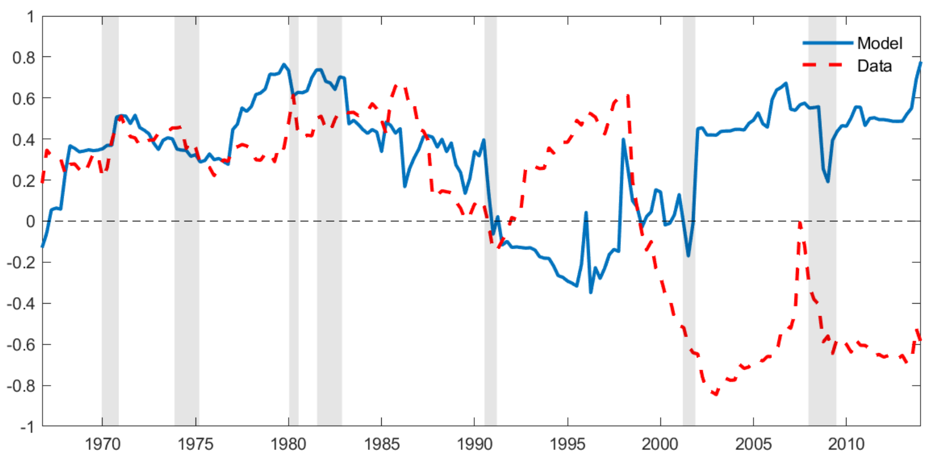

The figure depicts conditional 20-quarter rolling window correlations of stock returns and returns of 5-year nominal bonds. The red dashed line is based on the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve. The blue solid line is based on the returns which a constrained version of our model with CRRA preferences (i.e. $\gamma = \frac{1}{\psi}$) would have implied given the historical paths of consumption, inflation, and state variables. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is -0.15.

Figure 14: Dividend-price ratio with CRRA preferences

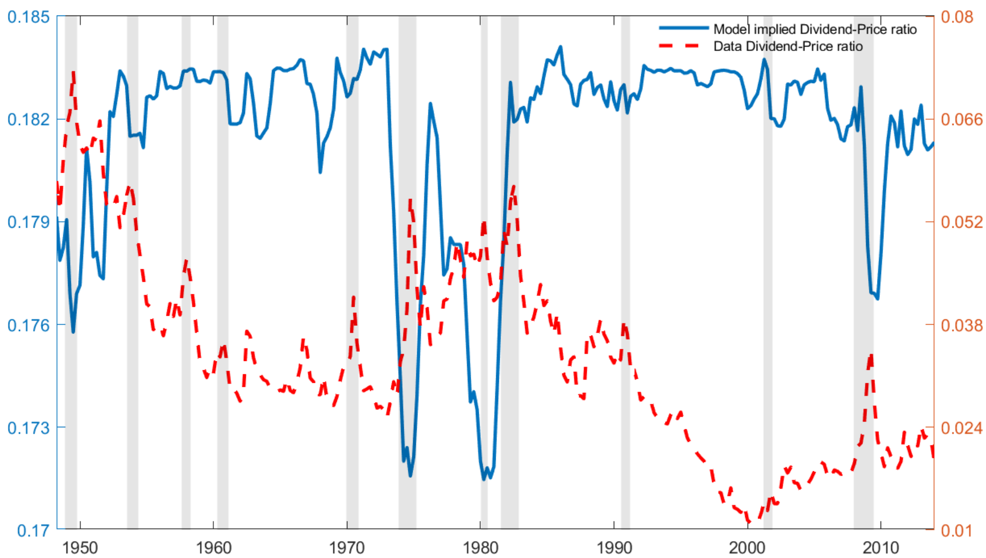

The figure depicts quarterly time series of dividend-price ratios. The blue solid line is the dividend-price ratio of the S&P 500 index. The red dashed line is obtained by plugging the historical paths of consumption, inflation, and our state variables $\widehat{p}_i$ into the numerical solution of the constrained model with CRRA preferences. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is -0.45.

Figure 15: Conditional stock-bond return correlation with full information

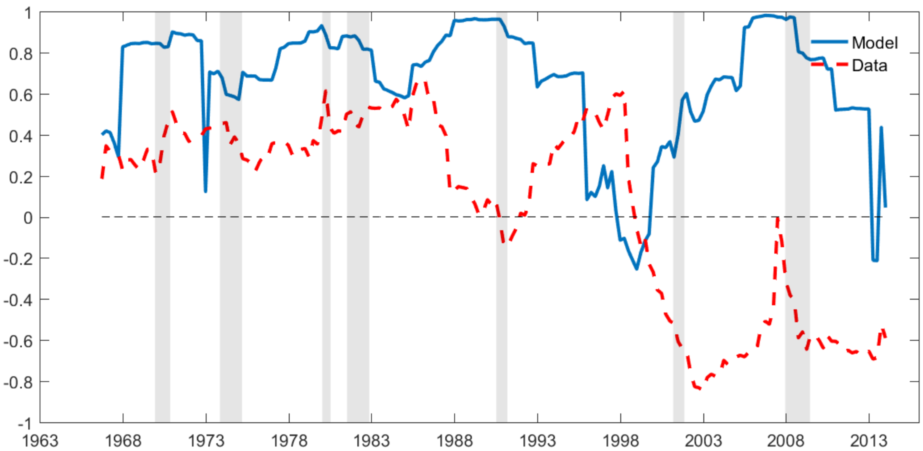

The figure depicts conditional 20-quarter rolling window correlations of stock returns and returns of 5-year nominal bonds. The red dashed line is based on the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve. The blue solid line is based on the returns which a version of our model with full information would have implied given the historical paths of consumption, inflation, and state variables. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is 0.09.

Figure 16: Time-varying probability of low expected consumption growth (robustness for different samples)

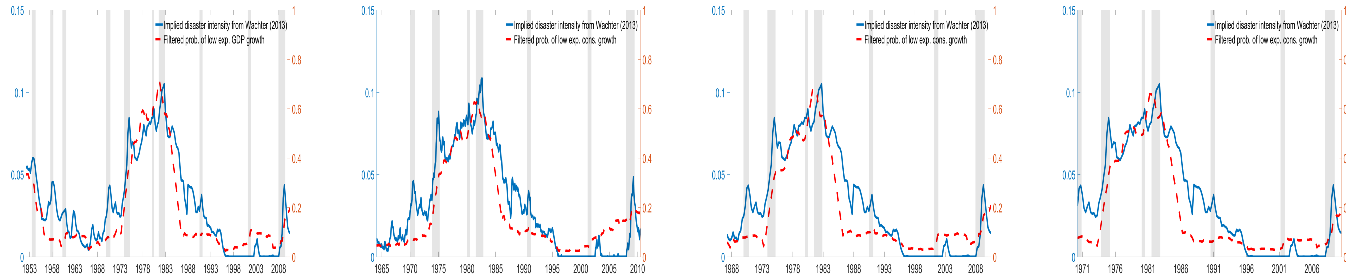

Estimation with inflation data

quarterly gdp growth (1947-2014)	monthly cons growth (1959-2014)	quarterly cons growth (1962-2014)	quarterly cons growth (1965-2014)

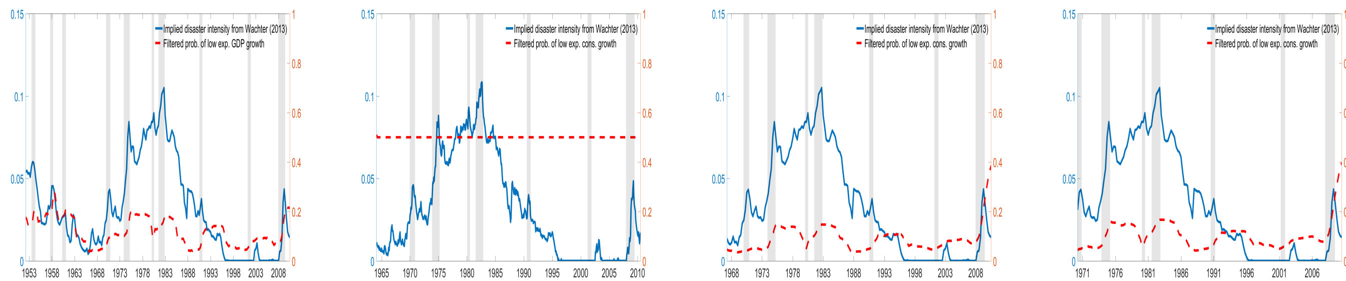

Estimation without inflation data

quarterly gdp growth (1947-2014)	monthly cons growth (1959-2014)	quarterly cons growth (1962-2014)	quarterly cons growth (1965-2014)

The figure depicts the results from applying the estimation methodology described in Section 3.2 to different samples. The upper graphs show the time series for the filtered probabilities for states with low expected consumption growth obtained using both consumption (or GDP) and inflation data. For the lower graphs we use consumption (or GDP) data only. The two pictures on the left hand side are based on quarterly GDP growth rates (1947Q1 to 2014Q1) instead of consumption growth rates. The next two pictures are based on an estimation using monthly consumption growth rates instead of quarterly consumption growth rates, which are available from 1959 onwards. The final four pictures are obtained from the subsamples of quarterly consumption growth rates starting in 1962 and 1965, respectively.

Figure 17: Time-varying probability of low expected consumption growth for constrained model (estimated from consumption growth and inflation)

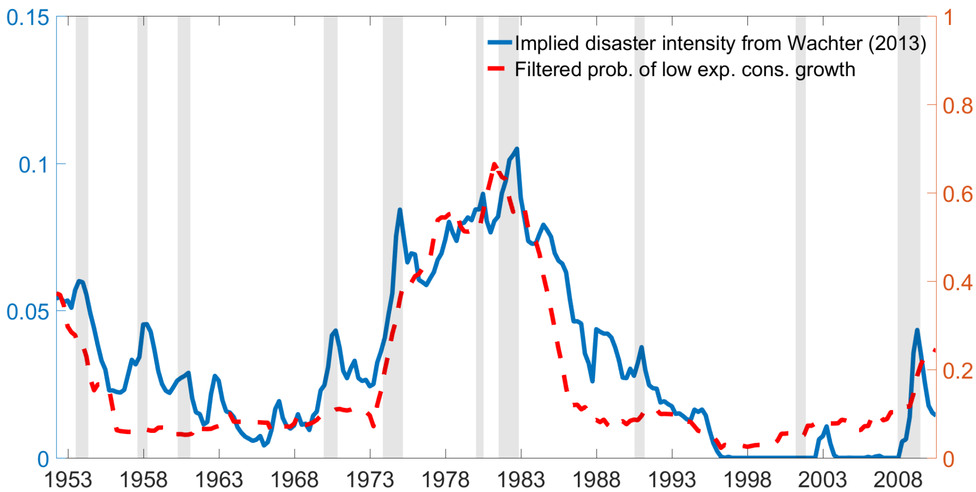

The blue solid line depicts the time-varying disaster intensity which Wachter (2013) extracts from asset price data. The red dashed line shows the 20-quarter moving average of the estimated $\widehat{p}_3 + \widehat{p}_4$ from the constrained model with equal expected consumption growth in states 1 and 2 and in states 3 and 4, respectively, for the period from 1947 to 2014. The correlation between the two time series is 0.88.

Table 1: Markov chain estimation

Panel A: Consumption and inflation parameters

	μ_1^C	μ_2^C	σ^C
Consumption growth	0.585	-0.450	0.421
	(0.094)	(0.301)	(0.047)

Panel B: Markov chain transition probabilities

	to state 1	to state 2
from state 1	0.939	0.061
	(0.092)	(0.092)
from state 2	0.467	0.533
	(0.217)	(0.217)

Panel C: Optimal number of states

	2 states	3 states
Log likelihood	-195.44	-184.14
Penalty term	39.14	72.68
Bayes Information Criterion	**430.02**	440.96
$(= -2 \cdot \log L + \text{penalty term})$		

Panel D: Wald test of $\mu_1^C = \mu_2^C$

Test statistic	p-value
14.818	0.000

This table reports the results from our univariate Markov chain estimation for consumption growth only. Growth rates are given in percentage points and quarterly. The data span the period from 1947 to 2014 at quarterly frequency. Numbers in parantheses are standard errors obtained from a standard block bootstrap with block length of 10 quarters. Panel D reports results from a Wald test of the null hypothesis $\mu_1^C = \mu_2^C$, based on the bootstrapped standard errors.

Table 2: Markov chain estimation

Panel A: Consumption and inflation quarterly parameters

	μ_1^i	μ_2^i	μ_3^i	μ_4^i	σ^i	ρ
Consumption growth	0.591	0.475	0.111	-0.249	0.504	-0.218
	(0.242)	(0.077)	(0.157)	(0.315)	(0.052)	(0.103)
Inflation	1.176	0.540	2.379	-0.729	0.309	
	(0.486)	(0.423)	(0.831)	(1.221)	(0.029)	

Panel B: Markov chain transition probabilities

	to state 1	to state 2	to state 3	to state 4
from state 1	0.909	0.027	0.039	0.025
	(0.202)	(0.087)	(0.160)	(0.049)
from state 2	0.022	0.970	0.008	0
	(0.029)	(0.056)	(0.027)	(0.024)
from state 3	0.135	0.037	0.828	0
	(0.064)	(0.077)	(0.088)	(0.028)
from state 4	0	0.337	0	0.663
	(0.065)	(0.253)	(0.152)	(0.203)

Panel C: Optimal number of states

	3 states	4 states	5 states	6 states
Log likelihood	-357.80	-317.65	-307.60	-287.75
Penalty term	119.40	169.70	238.80	320.50
Bayes Information Criterion	835.00	**805.00**	854.00	896.00
$(= -2 \cdot \log L + \text{penalty term})$				

Panel D: Wald test of $\mu_1^C = \mu_2^C = \mu_3^C = \mu_4^C$

Test statistic	p-value
12.051	0.007

This table reports the results from our bivariate Markov chain estimation for consumption growth and inflation. Growth rates are given in percentage points and quarterly. The data span the period from 1947 to 2014 at the quarterly frequency. Numbers in parentheses are standard errors obtained from a standard block bootstrap with block length of 10 quarters. Panel D reports results from a Wald test of the null hypothesis $\mu_1^C = \mu_2^C = \mu_3^C = \mu_4^C$, based on the bootstrapped standard errors.

Table 3: Unconditional moments of consumption growth and inflation

	Data		Model	
	Cons. growth	Inflation	Cons. growth	Inflation
Mean	0.456	0.868	0.453	0.828
			[0.403, 0.504]	[0.645, 1.017]
Volatility	0.5324	0.682	0.531	0.649
			[[0.498, 0.565]	[0.509, 0.786]
Correlation	-0.138		-0.100	
			[-0.241, 0.041]	
AC(1)	0.261	0.756	0.0695	0.622
			[-0.020, 0.161]	[0.474, 0.742]
AC(2)	0.237	0.631	0.049	0.520
			[-0.038, 0.135]	[0.362, 0.657]
AC(3)	0.208	0.578	0.031	0.439
			[-0.051, 0.114]	[0.271, 0.592]

The table reports unconditional moments of quarterly consumption growth and inflation. The columns labeled "Data" are based on quarterly data from 1947 to 2014. The columns labeled "Model" have been obtained by Monte Carlo simulation using the parameters given in Table 2 (5,000 paths of 68 years each). The numbers in parentheses give 90% confidence bounds around the point estimates.

Table 4: Regressions of stock return volatilities on state variables

Panel A: Model				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3 + \widehat{p}_4)$	Adj. R^2
0.127	0.014	0.010		0.811
(28.215)	(9.201)	(9.816)		
[0.084, 0.178]	[0.002, 0.018]	[0.000, 0.023]		[0.155, 0.761]
0.098			0.020	0.817
(24.194)			(11.388)	
[0.075, 0.133]			[0.004, 0.027]	[0.049, 0.721]

Panel B: Data				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3 + \widehat{p}_4)$	Adj. R^2
0.278	0.012	0.023		0.304
(12.679)	(1.416)	(4.475)		
0.214			0.022	0.214
(9.056)			(1.976)	

The table reports results from time series regressions. The dependent variables in each regression are volatilities of quarterly stock returns computed over rolling windows of 20 quarters. The independent variables are logarithms of the averages of the $\widehat{p}_i$ over the same 20 quarters periods. The regressions labeled "Data" are based on the estimated time series of the $\widehat{p}_i$ depicted in Figure 1 as well as the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve (see Gürkaynak, Sack, and Wright (2007)). The regressions labeled "Model" are based on the same time series of the $\widehat{p}_i$, but use the returns which our model would have implied given this path of consumption, inflation, and state variables. The financial data for these regressions starts in 1965. The model parameters are estimated using macroeconomic data since 1947. The numbers in parentheses are t-statistics adjusted following Newey and West (1987) (20 lags). The numbers in brackets denote 90% confidence bounds around the regression coefficients and are obtained from a Monte Carlo simulation of the model (5,000 paths of 68 years each, with the last 50 years of each path used in the regressions).

Table 5: Regressions of stock return volatilities on alternative explanatory variables

Panel A: Model				
const.	extreme entropy	expected inflation	realized inflation	Adj. R^2
0.005	0.291			0.460
(0.569)	(5.608)			
$[-0.019, 0.065]$	[0.020, 0.561]			[-0.002, 0.566]
0.031		0.007		0.427
(2.991)		(3.705)		
$[0.032, 0.078]$		[-0.009, 0.042]		[-0.009, 0.539]
0.024			0.008	0.443
(2.109)			(3.915)	
$[0.024, 0.079]$			[-0.009, 0.053]	[-0.009, 0.546]
Panel B: Data				
const.	extreme entropy	expected inflation	realized inflation	Adj. R^2
0.097	0.399			0.187
(2.607)	(2.320)			
0.147		0.006		0.036
(4.814)		(0.855)		
0.153			0.004	0.040
(6.010)			(0.781)	

The table reports results from time series regressions. The dependent variables in each regression are volatilities of quarterly stock returns computed over rolling windows of 20 quarters. The independent variables are the average of the extreme entropy ($\widehat{p}_3 \ln \widehat{p}_3 + \widehat{p}_4 \ln \widehat{p}_4$), the average expected inflation and the average realized inflation, always taken over the same 20 quarter periods. "Data" and "Model" have the same meaning as in Table 4.

Table 6: Regressions of stock-bond return correlations on state variables

Panel A: Model				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3 + \widehat{p}_4)$	Adj. R^2
0.282	0.304	-0.178		0.788
(3.315)	(11.599)	(-7.754)		
$[-7.465, 6.482]$	[0.079, 0.423]	[-0.515, -0.015]		[0.172, 0.725]
0.679			0.263	0.409
(4.193)			(3.830)	
$[-0.265, 0.995]$			[-0.111, 0.392]	[-0.008, 0.474]

Panel B: Data				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3 + \widehat{p}_4)$	Adj. R^2
-0.512	0.182	-0.294		0.363
(-1.61)	(2.731)	(-3.608)		
0.239			0.091	0.041
(0.923)			(0.774)	

The table reports results from time series regressions. The dependent variables in each regression are correlations of quarterly holding-period returns of stocks and 5-year nominal bonds computed over rolling windows of 20 quarters. The independent variables are logarithms of the averages of the $\widehat{p}_i$ over the same 20 quarters periods. "Data" and "Model" have the same meaning as in Table 4. Numbers in parentheses are t-statistics adjusted following Newey and West (1987) (20 lags), numbers in brackets are 90% confidence bounds around the regression coefficients and have been obtained from a Monte Carlo simulation of the model (5,000 paths of 68 years each, with the last 50 years of each path used in the regressions).

Table 7: Regressions of stock-bond return correlations on alternative explanatory variables

Panel A: Model				
const.	extreme entropy	expected inflation	realized inflation	Adj. R^2
-0.312	2.513			0.217
(-1.481)	(1.783)			
$[-1.315, 0.718]$	[-3.856, 8.848]			[-0.009, 0.313]
-0.473		0.154		0.645
(-5.036)		(10.846)		
$[-0.759, 0.069]$		[0.118, 1.024]		[0.005, 0.577]
-0.627			0.194	0.678
(-6.132)			(11.430)	
$[-0.928, 0.059]$			[0.117, 1.239]	[0.002, 0.558]
Panel B: Data				
const.	extreme entropy	expected inflation	realized inflation	Adj. R^2
-0.047	0.560			0.012
(-0.096)	(0.220)			
-0.613		0.167		0.334
(-2.028)		(3.172)		
-0.488			0.135	0.328
(-1.904)			(3.324)	

The table reports results from time series regressions. The dependent variables in each regression are correlations of quarterly holding-period returns of stocks and 5-year nominal bonds computed over rolling windows of 20 quarters. The independent variables are the average of the extreme entropy ($\widehat{p}_3 \ln \widehat{p}_3 + \widehat{p}_4 \ln \widehat{p}_4$), the average expected inflation and the average realized inflation, always taken over the same 20 quarter periods. "Data" and "Model" have the same meaning as in Table 4. Numbers in parentheses are t-statistics adjusted following Newey and West (1987) (20 lags), numbers in brackets are 90% confidence bounds around the regression coefficients and have been obtained from a Monte Carlo simulation of the model (5,000 paths of 68 years each, with the last 50 years of each path used in the regressions).

Table 8: Regressions of return correlations on state variables with CRRA preferences

Panel A: Model with CRRA preferences				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3 + \widehat{p}_4)$	Adj. R^2
0.996	0.143	0.092		0.756
(11.953)	(4.081)	(3.597)		
[0.338, 1.915]	[0.001, 0.264]	[-0.061, 0.355]		[0.056, 0.618]
0.719			0.196	0.753
(12.465)			(5.296)	
[0.388, 1.198]			[0.051, 0.396]	[0.016, 0.589]

Panel B: Data				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3 + \widehat{p}_4)$	Adj. R^2
-0.512	0.182	-0.294		0.363
(-1.61)	(2.731)	(-3.608)		
0.239			0.091	0.041
(0.923)			(0.774)	

The table reports results from time series regressions. The dependent variables in each regression are correlations of quarterly holding-period returns of stocks and 5-year nominal bonds computed over rolling windows of 20 quarters. The independent variables are logarithms of the averages of the $\widehat{p}_i$ over the same 20 quarters periods. "Data" and "Model" have the same meaning as in Table 4. Numbers in parentheses are t-statistics adjusted following Newey and West (1987) (20 lags). The numbers in brackets denote 90% confidence bounds around the regression coefficients and are obtained from a Monte Carlo simulation of the model (5,000 paths of 68 years each, with the last 50 years of each path used in the regressions).

Table 9: Unconditional asset pricing moments

		Data	Benchmark	Param. 2	Param. 3	Param. 4	Combined
nominal equity return	mean	0.116	0.074	0.079	0.114	0.081	0.127
			(0.009)	(0.008)	(0.010)	(0.013)	(0.013)
	vol	0.148	0.075	0.079	0.117	0.123	0.121
			(0.004)	(0.004)	(0.009)	(0.007)	(0.009)
	AC(1)	0.079	-0.054	-0.075	-0.131	-0.062	-0.141
			(0.044)	(0.043)	(0.050)	(0.043)	(0.050)
nominal 3-month risk-free rate	mean	0.043	0.062	0.063	0.066	0.062	0.066
			(0.005)	(0.004)	(0.008)	(0.005)	(0.006)
	vol	0.050	0.006	0.006	0.005	0.006	0.010
			(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
	AC(1)	0.987	0.774	0.752	0.620	0.775	0.632
			(0.065)	(0.066)	(0.101)	(0.065)	(0.099)
real dividend growth	mean	0.067	0.018	0.018	0.016	0.018	0.016
			(0.008)	(0.006)	(0.006)	(0.008)	(0.006)
	vol	0.139	0.051	0.051	0.052	0.082	0.052
			(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
	AC(1)	-0.251	0.030	0.031	0.055	0.030	0.053
			(0.038)	(0.038)	(0.042)	(0.038)	(0.042)
yield spread (5y-3m)	mean	0.010	0.001	0.001	-0.012	0.001	-0.024
			(0.002)	(0.002)	(0.003)	(0.002)	(0.003)
	vol	0.012	0.011	0.011	0.012	0.011	0.014
			(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
uncond. stock-bond corr.		0.114	0.105	0.083	-0.316	0.188	-0.194
			(0.128)	(0.130)	(0.151)	(0.110)	(0.153)

The table shows unconditional asset pricing moments. "Data" refers to the CRSP value-weighted index for stocks and to the data set provided by Gürkaynak, Sack, and Wright (2007) for bonds. The model-implied values are computed via Monte Carlo simulation. All numbers are computed based on monthly observations and then annualized. In the data the average yield spread (5y - 3m) is available from 1952 onwards. The correlation between nominal stock and 5y-bond returns is based on five-year rolling window estimates with data from 1962 onwards. The column labeled Benchmark presents results for the model with the estimated parameters from Table 2. Parametrization 2 is the same as the benchmark parametrization except that we lower the expected consumption growth in states 3 and 4 by one standard error, i.e. we set $\mu_3^C = 0.129$ and $\mu_4^C = -1.628$. Parametrization 3 is the same as the benchmark parametrization except that we increase the transition probabilities from states 3 and 4 to the other states by one standard error, i.e. we set the third and fourth row of the transition matrix to (0.057 0.027 0.916 0.000) and (0.000 0.137 0.000 0.863), respectively. Parametrization 4 is the same as the benchmark parametrization except for a leverage parameter $\phi = 8$. The column labeled "Combined" reports results when Parametrization 2 and 3 are applied jointly.

ONLINE APPENDIX

Extreme Inflation
and Time-Varying Expected Consumption Growth

Ilya Dergunov* Christoph Meinerding** Christian Schlag*

First version: July 13, 2016
This version: January 3, 2022

Abstract

This Online Appendix serves as a companion to our paper "Extreme Inflation and Time-Varying Expected Consumption Growth". It provides additional results and explanations not reported in the main text due to space constraints.

*Faculty of Economics and Business Administration and Research Center SAFE, Goethe University, 60323 Frankfurt am Main, Germany. E-mails: dergunov@safe.uni-frankfurt.de, schlag@finance.uni-frankfurt.de.

**Deutsche Bundesbank, Research Centre, 60431 Frankfurt am Main, Germany. E-mail: christoph.meinerding@bundesbank.de.

A. Maximum Likelihood Estimation

The maximum likelihood estimation is based on Hamilton (1994). The parameters of the model and the transitions probabilities p_{ij} are collected in the vector Θ. This vector Θ is estimated based on the data Y_T that is observed until time T. This data forms a $T \times 2$ matrix. Let $P(s_t = j|\Theta, Y_t)$ define the probability of state s_t conditional on all data until t and conditional on the knowledge of the parameters Θ. The econometrician assigns $P(s_t = j|\Theta, Y_t)$ to the possibility that the observation at time t is generated by state j. These conditional probabilities are collected in the $n \times 1$ vector $\hat{\xi}_{t|t}$, where n denotes the number of states, which is fixed throughout the entire estimation. We have

$$\begin{aligned} \hat{\xi}_{t|t} &= \frac{\hat{\xi}_{t|t-1} \odot \eta_t}{\mathbf{1'} \cdot (\hat{\xi}_{\mathbf{t|t-1}} \odot \eta_{\mathbf{t}})} \\ \hat{\xi}_{t+1|t} &= \mathbf{Q} \cdot \hat{\xi}_{\mathbf{t|t}}. \end{aligned}$$

Here η_t is the conditional density, whose jth element is

$$f(y_t|s_t = j; \Theta) = \frac{1}{\sqrt{2\pi|\Omega|}} \exp\left\{0.5 \cdot (y_t - \mu_j)'\Omega^{-1}(y_t - \mu_j)\right\},$$

where y_t is an $n \times 1$ vector, $\mu_j = (\mu_j^C, \mu_j^\pi)'$ and $\Omega = \Sigma\Sigma'$ where

$$\Sigma = \begin{pmatrix} \sigma^C\sqrt{1-\rho^2} & \sigma^C\rho \\ 0 & \sigma^\pi \end{pmatrix}$$

$\mathbf{Q}$ is the transition probability matrix

$$\mathbf{Q} = \begin{pmatrix} q_{11} & \dots & q_{1n} \\ \vdots & \ddots & \vdots \\ q_{n1} & \dots & q_{nn} \end{pmatrix},$$

$\odot$ denotes element-by-element multiplication, and $\mathbf{1}$ denotes a vector of ones. The jth element of the product $\hat{\xi}_{t|t-1} \odot \eta_t$ is interpreted as the conditional joint density of $\mathbf{y_t}$ and s_t

$$P(s_t = j|\Theta) \times f(y_t|s_t = j; \Theta) = p(\mathbf{y_t}, \mathbf{s_t} = \mathbf{j}|\mathbf{\Theta}).$$

The density of the observed vector $\mathbf{y_t}$ is the sum

$$f(y_t|\Theta) = \mathbf{1'} \cdot (\hat{\xi}_{\mathbf{t|t-1}} \odot \eta_{\mathbf{t}})$$

The objective is to find a maximum of the log likelihood function

$$\mathbb{L}(\Theta) = \sum_{t=1}^{T} \log f(y_t|\Theta).$$

The starting value for the maximization, $\hat{\xi}_{1|0}$, is set to $\frac{1}{n} \cdot \mathbf{1}$.

B. Filtering

B.1. Dynamics of the state variables

The dynamics of consumption and the log price level can be rewritten as

$$\begin{aligned} d\ln C_t &= \left(\sum_{i=1}^{n} \mu_i^C \mathbb{1}_{\{S_t=i\}}\right) dt + \sigma^C \left(\sqrt{1-\rho^2} dW_t^C + \rho dW_t^\pi\right) \\ d\pi_t &= \sum_{i=1}^{n} \mu_i^\pi \mathbb{1}_{\{S_t=i\}} + \sigma^\pi dW_t^\pi, \end{aligned}$$

where $\mathbb{1}_{\{S_t=i\}}$ is the indicator function equal to one, if the economy is in state i at time t, $(i = 1, \ldots, n)$, and equal to zero otherwise. In matrix form, this becomes

$$\begin{pmatrix} d\ln C_t \\ d\pi_t \end{pmatrix} = \begin{pmatrix} \sum_{i=1}^{n} \mu_i^C \mathbb{1}_{\{S_t=i\}} \\ \sum_{i=1}^{n} \mu_i^\pi \mathbb{1}_{\{S_t=i\}} \end{pmatrix} dt + \Sigma \begin{pmatrix} dW_t^C \\ dW_t^\pi \end{pmatrix},$$

with

$$\Sigma = \begin{pmatrix} \sigma^C \sqrt{1-\rho^2} & \sigma^C \rho \\ 0 & \sigma^\pi \end{pmatrix}$$

and $d[W^C, W^\pi] = 0$.

The inverse of the of the volatility matrix Σ is

$$\Sigma^{-1} = \frac{1}{\sigma^C \sigma^\pi \sqrt{1-\rho^2}} \begin{pmatrix} \sigma^\pi & -\sigma^C \rho \\ 0 & \sigma^C \sqrt{1-\rho^2} \end{pmatrix}.$$

We assume that the drift rates are unobservable and need to be filtered by the investor. Mathematically, there are two filtrations, $\mathcal{F}$ and $\mathcal{G}$, where $\mathcal{F}$ is generated by the processes $(C_t)_t$, $(\pi_t)_t$ and $(S_t)_t$, whereas $\mathcal{G} \subset \mathcal{F}$ is generated by the processes $(C_t)_t$ and $(\pi_t)_t$ only. The equilibrium in the economy is based on the dynamics of $(C_t)_t$ and $(\pi_t)_t$ under the investor filtration $\mathcal{G}$, i.e. on the projections $\widehat{\mu}_t^C = E\left[\mu^C(S_t)|\mathcal{G}_t\right] = \sum_{i=1}^n \widehat{p}_{it}\mu_i^C$ and $\widehat{\mu}_t^\pi = E\left[\mu^\pi(S_t)|\mathcal{G}_t\right] = \sum_{i=1}^n \widehat{p}_{it}\mu_i^\pi$, where $\widehat{p}_{it} = E\left[\mathbb{1}_{\{S_t=i\}}|\mathcal{G}_t\right]$.

An application of Theorem 9.1 on p. 355 of Liptser and Shiryaev (2001) yields dynamics for the projected (henceforth also called "subjective") probabilities:

$$d\widehat{p}_{it} \quad = \quad \left(\widehat{p}_{it}\lambda_{ii} + \sum_{j\neq i}\widehat{p}_{jt}\lambda_{ji}\right)dt \quad + \quad \widehat{p}_{it}\left[\begin{pmatrix}\mu_i^C \\ \mu_i^\pi\end{pmatrix} - \sum_{j=1}^n \widehat{p}_{jt}\begin{pmatrix}\mu_j^C \\ \mu_j^\pi\end{pmatrix}\right]' \times \Sigma'^{-1}\begin{pmatrix}d\widehat{W}_t^C \\ d\widehat{W}_t^\pi\end{pmatrix},$$

where

$$\begin{pmatrix}d\widehat{W}_t^C \\ d\widehat{W}_t^\pi\end{pmatrix} = \Sigma^{-1}\left[\begin{pmatrix}\mu_i^C \\ \mu_i^\pi\end{pmatrix} - \sum_{j=1}^n \widehat{p}_{jt}\begin{pmatrix}\mu_j^C \\ \mu_j^\pi\end{pmatrix}\right]dt + \begin{pmatrix}dW_t^C \\ dW_t^\pi\end{pmatrix}.$$

In particular, $d[\widehat{W}^C, \widehat{W}^\pi] = 0$. Under the investor's filtration, log consumption dynamics are given as

$$d\ln C_t = \sum_{i=1}^n \mu_i^C \widehat{p}_{it}dt + \sigma^C\left(\sqrt{1-\rho^2}d\widehat{W}^C + \rho d\widehat{W}^\pi\right)$$

For notational convenience, we define

$$\sigma_{\widehat{p}_i} d\widehat{W}_t \equiv \widehat{p}_{it}\left[\begin{pmatrix}\mu_i^C \\ \mu_i^\pi\end{pmatrix} - \sum_{j=1}^n \widehat{p}_{jt}\begin{pmatrix}\mu_j^C \\ \mu_j^\pi\end{pmatrix}\right]' \times \Sigma'^{-1} \times \begin{pmatrix}d\widehat{W}_t^C \\ d\widehat{W}_t^\pi\end{pmatrix}$$

and

$$\begin{aligned}
\sigma_{c,\hat{p}_i} &\equiv \left(\rho\sigma^C, \quad \sqrt{1-\rho^2}\sigma^C\right) \times \sigma'_{\hat{p}_i} \\
\sigma_{\pi,\hat{p}_i} &\equiv (0, \quad \sigma^\pi) \times \sigma'_{\hat{p}_i}.
\end{aligned}$$

B.2. Discussion

In order to understand the dynamics of $\widehat{p}_{it}$, it is instructive to analyze the drift and the diffusion component separately. The drift in (4) is a linear function of the transition intensities λ and the current estimates of the probabilities $\widehat{p}$. Since the states (and consequently also switches between states) are unobservable, the subjective probability of being in state i changes deterministically over time, depending on the conditional probabilities to enter or exit state i. The drift therefore comprises two terms. The first term, $\widehat{p}_{it}\lambda_{ii} = -\widehat{p}_{it}\sum_{j\neq i}\lambda_{ij}$, involves the intensities for a switch from state i to some other state $j \neq i$. Loosely speaking, the more time goes by, the higher the chance that an unobserved switch from state i to some other state j has occurred in the meantime. This effect induces a negative drift for $\widehat{p}_{it}$, i.e., the estimated probability of still being in state i decreases in expectation. The second term, $\sum_{j\neq i}\widehat{p}_{jt}\lambda_{ji}$, captures the probabilities of entering state i, given that the economy is currently in a state j different from i. Suppose one of the $\widehat{p}_{jt}$ $(j \neq i)$ is currently large. Then, as time passes and if no other conflicting signals arrive, it becomes more and more likely that an unobserved switch to state i has occurred in the meantime. This effect induces a positive drift in $\widehat{p}_{it}$. The overall sign of the drift of $\widehat{p}_{it}$ thus depends on the current estimate of all state probabilities. In particular, the drift terms ensure that the probabilities $\widehat{p}$ will always be between zero and one.

The volatility of the change in $\widehat{p}_i$ is a *quadratic* function of all probabilities $\widehat{p}_j$ $(j = 1, \ldots, n)$. The probability update is largest when the investor is rather uncertain about the current state of the economy, i.e., for intermediate values of $\widehat{p}_i$. When the investor is almost sure about the current state of the economy (i.e., when one of the $\widehat{p}_j$ is close to one and the others are close to zero), the estimated probability will change in a basically deterministic fashion. To see this, note that when the respective estimate $\widehat{p}_i$ is close to zero, the diffusion term in (4) is obviously also close to zero, since $\widehat{p}_i$ is one factor of the product in front of the Wiener innovations. When $\widehat{p}_i$ is in turn close to one, the term in square brackets in (4) will be very close to zero, since in the sum only the term involving $\widehat{p}_i$ will remain, whereas all the others will vanish.

The volatility of the innovation in the filtered probability also depends on the precision of the signals. When the signals are very imprecise, i.e., when the volatilities σ^C and σ^π are large, an observed innovation in (log) consumption growth or inflation delivers less information about the true state, and the investor will put less weight on them when computing the new estimate for p_i.

The sign of the diffusion term depends on the sign of the 'observed' Brownian shocks $d\widehat{W}$. These are defined via the restriction that $\ln C$ and π are observable, so technically speaking they have to be adapted to both $\mathcal{F}$ and $\mathcal{G}$, which implies $\mu(S_t)dt + \Sigma dW_t = \widehat{\mu}_t dt + \Sigma d\widehat{W}_t$.

C. Solution of the Model with Partial Information

C.1. Wealth-consumption ratio

The indirect utility function of the investor is given by

$$J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}) = E_t \left[\int_t^\infty f(C_s, J(C_s, \widehat{p}_{1s}, \ldots, \widehat{p}_{ns}))\, ds \right].$$

$J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}) + \int_0^t f(C_s, J(C_s, \widehat{p}_{1s}, \ldots, \widehat{p}_{ns}))\, ds$ is a martingale, therefore we have the Bellman equation

$$E[dJ(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}) + f(C_t, J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}))\, dt] = 0,$$

or, equivalently,

$$\frac{\mathcal{A}J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt})}{J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt})} + \frac{f(C_t, J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}))}{J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt})} = 0, \tag{C.1}$$

where $\mathcal{A}$ is the infinitesimal generator. The aggregator function is given by

$$f(C, J) = \frac{\beta C^{1-\frac{1}{\psi}}}{\left(1 - \frac{1}{\psi}\right) \left[(1-\gamma)J\right]^{\frac{1}{\theta}-1}} - \beta\theta J.$$

We conjecture a functional form for J:

$$J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}) \quad = \quad \frac{C_t^{1-\gamma}}{1-\gamma} \left(\beta e^{v(\widehat{p}_{1t}, \ldots, \widehat{p}_{nt})}\right)^\theta$$

where in the end v will turn out to be the log wealth-consumption ratio. This functional form together with the definition of $f(C, J)$ implies

$$\frac{f(C,J)}{J} = \theta e^{-v} - \theta\beta.$$

With $I \equiv e^v$ the partial derivatives of J (denoted by subscripts) are

$$\begin{aligned}
J_c &= C^{-\gamma}\left(\beta e^v\right)^{\theta} \\
J_{cc} &= -\gamma C^{-\gamma-1}\left(\beta e^v\right)^{\theta} \\
J_{\widehat{p}_i} &= \frac{C^{1-\gamma}}{1-\gamma}\beta^{\theta}\theta I^{\theta-1} I_{\widehat{p}_i} \\
J_{\widehat{p}_i\widehat{p}_j} &= \frac{C^{1-\gamma}}{1-\gamma}\beta^{\theta}\theta\left[(\theta-1)I^{\theta-2}I_{\widehat{p}_i}^2 + I_{\widehat{p}_i\widehat{p}_j} I^{\theta-1}\right] \\
J_{c\widehat{p}_i} &= C^{-\gamma}\beta^{\theta}\theta I^{\theta-1} I_{\widehat{p}_i},
\end{aligned}$$

resulting in

$$\begin{aligned}
\frac{J_c}{J}dC_t &= (1-\gamma)\sum_{i=1}^{n}\mu_i^C\widehat{p}_{it}dt + \frac{1}{2}(\sigma^C)^2(1-\gamma)dt + (1-\gamma)\sigma^C\left(\sqrt{1-\rho^2}d\widehat{W}_t^C + \rho d\widehat{W}_t^{\pi}\right) \\
\frac{J_{cc}d[C,C]_t}{J} &= (\sigma^C)^2(-\gamma)(1-\gamma)dt \\
\frac{J_{c\widehat{p}_i}d[C,\widehat{p}_i]_t}{J} &= \theta(1-\gamma)\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i}dt \\
\frac{J_{\widehat{p}_i}d\widehat{p}_{it}}{J} &= \theta\frac{I_{\widehat{p}_i}}{I}d\widehat{p}_{it} \\
\frac{J_{\widehat{p}_i\widehat{p}_j}d[\widehat{p}_i,\widehat{p}_j]_t}{J} &= \frac{1}{2}\theta d[\widehat{p}_i,\widehat{p}_j]_t\left[(\theta-1)\left(\frac{I_{\widehat{p}_i}}{I}\right)^2 + \frac{I_{\widehat{p}_i\widehat{p}_j}}{I}\right].
\end{aligned}$$

Plugging everything into (C.1) results in the following partial differential equation for I:

$$\begin{aligned}
0 = {} & \left[(1-\gamma)\sum_{i=1}^{n}\mu_i^C\widehat{p}_{it} + \frac{1}{2}(1-\gamma)^2(\sigma^C)^2 - \beta\theta\right] + \theta I^{-1} \\
& + \sum_{i=1}^{n-1}\theta\frac{I_{\widehat{p}_{it}}}{I}\left(\widehat{p}_i\lambda_{ii} + \sum_{\substack{j=1 \\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right) + \sum_{i=1}^{n-1}\theta(1-\gamma)\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i} \\
& + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\theta\left[(\theta-1)\left(\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I}\right) + \frac{I_{\widehat{p}_i\widehat{p}_j}}{I}\right]\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j},
\end{aligned} \tag{C.2}$$

There are $n-1$ state variables due to the restriction $\sum_{i=1}^{n} \widehat{p}_{it} = 1$. We solve the PDE with a Chebyshev approximation similar to Benzoni, Collin-Dufresne, and Goldstein (2011). We guess the following functional form for I as a function of the vector $\widehat{p} = (\widehat{p}_1, \widehat{p}_2, ..., \widehat{p}_{n-1})$ with:

$$\begin{aligned} I(\widehat{p}) &= \exp(B(\widehat{p})) \\ B(\widehat{p}) &= \sum_{j=0}^{d} \alpha_j T_j(\widehat{p}), \end{aligned} \tag{C.3}$$

where the $T_j(\widehat{p})$ are multivariate Chebyshev polynomials. For the interval $[-1,1]$, the univariate Chebyshev polynomials are defined recursively via

$$\begin{aligned} T_0(x) &= 1 \\ T_1(x) &= x \\ T_{d+1}(x) &= 2xT_d(x) - T_{d-1}(x). \end{aligned}$$

Univariate Chebyshev polynomials for the general interval $[a,b]$ are given by transformations

$$T_d\left(\frac{2x-b-a}{b-a}\right).$$

Multivariate versions of the Chebyshev polynomials are defined as sums of products of the univariate ones. The derivatives of the guess in (C.3) are

$$\begin{aligned} I_{\widehat{p}_i} &= e^{B(\widehat{p})} B_{\widehat{p}_i} = e^{B(\widehat{p})} \sum_{j=1}^{d} \alpha_j \frac{\partial T_j}{\partial \widehat{p}_i}(\widehat{p}) \\ I_{\widehat{p}_i \widehat{p}_j} &= e^{B(\widehat{p})} \left[(B_{\widehat{p}_i})^2 + B_{\widehat{p}_i \widehat{p}_j} \right] = e^{B(\widehat{p})} \left[\left(\sum_{j=1}^{d} \alpha_j \frac{\partial T_j}{\partial \widehat{p}_i}(\widehat{p}) \right)^2 + \left(\sum_{j=2}^{d} \alpha_j \frac{\partial^2 T_j}{\partial \widehat{p}_i \partial \widehat{p}_j}(\widehat{p}) \right) \right]. \end{aligned}$$

Plugging the partial derivatives into (C.2) gives

$$
\begin{aligned}
0 \;=\; & \left[(1-\gamma)\sum_{i=1}^{n}\mu_i^C \widehat{p}_i + \frac{1}{2}(1-\gamma)^2(\sigma^C)^2 - \beta\theta\right] \quad + \quad e^{-\sum_{j=0}^{d}\alpha_j T_j(\widehat{p})}\theta \\
& + \sum_{i=1}^{n-1}\theta\left(\sum_{j=1}^{d}\alpha_j \frac{\partial T_j}{\partial \widehat{p}_i}(\widehat{p})\right)\left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right) \quad + \quad \sum_{i=1}^{n-1}\theta(1-\gamma)\left(\sum_{j=1}^{d}\alpha_j\frac{\partial T_j}{\partial \widehat{p}_i}(\widehat{p})\right)\sigma_{c,\hat{p}_i} \\
& + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{k=1}^{n-1}\theta\left[(\theta-1)\sum_{j=1}^{d}\alpha_j\frac{\partial T_j}{\partial \widehat{p}_i}(\widehat{p})\sum_{j=1}^{d}\alpha_j\frac{\partial T_j}{\partial \widehat{p}_k}(\widehat{p}) + \sum_{j=2}^{d}\alpha_j\frac{\partial^2 T_j}{\partial \widehat{p}_i \partial \widehat{p}_k}(\widehat{p})\right]\sigma_{\hat{p}_i}\sigma'_{\hat{p}_k}.
\end{aligned}
$$

This equation is defined on the simplex Δ^{n-1}. We partition this simplex by choosing grid points according to the Chebyshev methodology. Evaluating the equation on every grid point leaves us with a number of algebraic equations, whose solution gives the Chebyshev coefficients α_j. For the multivariate Chebyshev polynomials we choose the order of $d = 4$. We have also tried higher values for d but the solution remained unchanged.

C.2. Pricing kernel

The pricing kernel is given by

$$
\xi_t \;=\; \exp\left(\int_0^t -\beta\theta - (1-\theta)I^{-1}(\widehat{p}_s)ds\right) C_t^{-\gamma}(I(\widehat{p}_t))^{\theta-1}
$$

with dynamics

$$
\begin{aligned}
\frac{d\xi_t}{\xi_t} \;=\; & -\beta\theta dt - (1-\theta)I^{-1}\,dt - \gamma\frac{dC_t}{C_t} + \frac{1}{2}\gamma^2(\sigma^C)^2 dt \\
& - (1-\theta)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}d\widehat{p}_{it} + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}(\theta-1)\left[\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + (\theta-2)\left(\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I^2}\right)\right]\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}dt \\
& - \gamma(\theta-1)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i}dt,
\end{aligned}
$$

where $\widehat{p} = (\widehat{p}_1, \widehat{p}_2, ..., \widehat{p}_n)$. For later use, we abbreviate the drift term as

$$\mu_{\xi,t} \equiv -\beta\theta \quad - (1-\theta)I^{-1}(\widehat{p}_t) \quad - \gamma\sum_{i=1}^{n}\mu_i^C\widehat{p}_{it} \quad + \frac{1}{2}\gamma^2(\sigma^C)^2 \quad - \sum_{i=1}^{n-1}(1-\theta)\frac{I_{\widehat{p}_i}}{I}\left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right)$$
$$+ \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}(\theta-1)\Big[\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + (\theta-2)\left(\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I^2}\right)\Big]\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j} \quad - \gamma(\theta-1)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i}.$$

C.3. Price-dividend ratio

We want to price a claim on levered consumption. Under the investor's filtration, dividends follow the process

$$d\ln D_t = \bar{\mu}dt + \phi\left(\sum_{i=1}^{n}(\mu_i^C - \bar{\mu})\widehat{p}_{it}\right)dt + \phi\sigma^C\left(\sqrt{1-\rho^2}d\widehat{W}_t^C + \rho d\widehat{W}_t^\pi\right).$$

Let ω denote the log price-dividend ratio. For $g(\xi, D, \omega) \equiv \xi De^\omega$, the Feynman-Kac formula yields

$$\frac{\mathcal{A}g(\xi, D, \omega)}{g(\xi, D, \omega)} + e^{-\omega} = 0. \tag{C.4}$$

Itô's Lemma gives

$$\frac{\mathcal{A}g_t}{g_t} = \mu_{\xi,t} + \mu_{D,t} + \mu_{\omega,t} + \frac{1}{2}\frac{d[\omega]_t}{dt} + \frac{d\,[\xi, D]_t}{\xi Ddt} + \frac{d\,[\omega, D]_t}{Ddt} + \frac{d\,[\omega, \xi]_t}{\xi dt}.$$

Another application of Itô's Lemma leads to

$$d\omega_t = \sum_{i=1}^{n-1}\omega_{\widehat{p}_i}d\hat{p}_{it} + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\omega_{\widehat{p}_i\widehat{p}_j}\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}dt$$

where the subscripts denote partial derivatives with respect to the state variables $\widehat{p}_i$. We can rewrite the drift of ω as a function of the derivatives $\omega_{\widehat{p}_i}$ and $\omega_{\widehat{p}_i,\widehat{p}_j}$ and the state variables $\widehat{p}_i$:

$$\mu_{\omega,t} = \sum_{i=1}^{n-1}\omega_{\widehat{p}_i}\left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right) \quad + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\omega_{\widehat{p}_i\widehat{p}_j}\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}.$$

The quadratic variation terms are:

$$
\begin{aligned}
d[\omega]_t &= \sum_{i=1}^{n-1}\sum_{j=1}^{n-1} \omega_{\widehat{p}_i}\omega_{\widehat{p}_j}\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}dt \\
\frac{d[\xi,\omega]_t}{\xi_t} &= -(1-\theta)\left(\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\omega_{\widehat{p}_i} + \sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\omega_{\widehat{p}_j}\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}dt - \gamma\sum_{i=1}^{n-1}\omega_{\widehat{p}_i}\sigma_{c,\hat{p}_i}dt \\
\frac{d[\xi,D]_t}{\xi_t D_t} &= -\gamma\phi(\sigma^C)^2 - (1-\theta)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\phi\sigma_{c,\hat{p}_i}dt \\
\frac{d[\omega,D]_t}{D_t} &= \sum_{i=1}^{n-1}\omega_{\widehat{p}_i}\phi\sigma_{c,\hat{p}_i}dt.
\end{aligned}
$$

Plugging everything into (C.4) gives the following PDE for ω:

$$
\begin{aligned}
0 \;=\; & -\beta\theta - (1-\theta)I^{-1} + e^{-\omega} - \gamma\sum_{i=1}^{n}\mu_i^C\widehat{p}_{it} + \bar{\mu} + \phi\left(\sum_{i=1}^{n}(\mu_i^C - \bar{\mu})\widehat{p}_{it}\right) + \frac{1}{2}(\phi-\gamma)^2(\sigma^C)^2 \\
& + \sum_{i=1}^{n-1}\left((\theta-1)\frac{I_{\widehat{p}_i}}{I} + \omega_{\widehat{p}_i}\right)\left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right) \quad + \sum_{i=1}^{n-1}\left((\phi-\gamma)(\theta-1)\frac{I_{\widehat{p}_i}}{I} + (\phi-\gamma)\omega_{\widehat{p}_i}\right)\sigma_{c,\hat{p}_i} \\
& + \sum_{i=1}^{n-1}\left(\frac{1}{2}(\theta-1)(\theta-2)\left(\frac{I_{\widehat{p}_i}}{I}\right)^2 + (\theta-1)\frac{I_{\widehat{p}_i}}{I}\omega_{\widehat{p}_i} + \frac{1}{2}\omega_{\widehat{p}_i}^2\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_i} \\
& + \sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\left(\frac{1}{2}(\theta-1)(\theta-2)\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I^2} + (\theta-1)\frac{I_{\widehat{p}_i}}{I}\omega_{\widehat{p}_j} + \frac{1}{2}\omega_{\widehat{p}_i}\omega_{\widehat{p}_j}\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j} \\
& + \frac{1}{2}\sum_{i=1}^{n-1}\left((\theta-1)\left(\frac{I_{\widehat{p}_i\widehat{p}_i}}{I}\right) + \omega_{\widehat{p}_i\widehat{p}_i}\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_i} \quad + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\left((\theta-1)\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + \omega_{\widehat{p}_i\widehat{p}_j}\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}.
\end{aligned}
$$

Similar to the wealth-consumption ratio, we approximate the price-dividend ratio $U \equiv e^{\omega}$ via a multivariate Chebyshev polynomial expansion, i.e., we approximate the function $U(\widehat{p})$ as

$$
U(\widehat{p}) \;=\; \exp\left\{\sum_{j=0}^{d}\beta_j T_j(\widehat{p})\right\},
$$

and solve the PDE numerically.

C.4. Pricing real bonds

Let the price of a real bond expiring at time T be denoted by $B_t^T = E_t[\frac{\xi_T}{\xi_t}]$ with the real pricing kernel

$$\frac{\xi_T}{\xi_t} = \beta^\theta \left(\frac{C_T}{C_t}\right)^{-\gamma} \exp\left\{-\beta\theta(T-t) + (\theta-1)\left(\int_t^T I^{-1}(\hat{p}_s)ds\right)\right\}(I(\hat{p}_t))^{\theta-1}$$

Let $b_t \equiv \ln B_t^T$ be the log real bond price (we will omit the superscript T in the following). The Feynman-Kac formula applied to $H(\xi_t, b_t) = \xi_t e^{b_t}$ yields the partial differential equation:

$$0 \;=\; \mathcal{A}H \;=\mu_{\xi,t} + \mu_{b,t} + \frac{1}{2}\frac{d[b]_t}{dt} + \frac{d[\xi,b]_t}{\xi_t dt}, \tag{C.5}$$

where μ_ξ and μ_b are the drifts of the processes ξ_t and b_t respectively. The dynamics of b_t are

$$db_t = \frac{\partial b_t}{\partial t}dt + \sum_{i=1}^{n-1} b_{\widehat{p}_i} d\widehat{p}_{it} + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1} b_{\widehat{p}_i\widehat{p}_j}\sigma_{\widehat{p}_i}\sigma'_{\widehat{p}_j}dt.$$

Plugging everything into (C.5) gives the following PDE for b_t:

$$\begin{aligned}
0 \;=\;& -\beta\theta - (1-\theta)I^{-1} - \gamma\sum_{i=1}^{n}\mu_i^C\widehat{p}_{it} + \frac{1}{2}\gamma^2(\sigma^C)^2 + \frac{\partial b_t}{\partial t} \\
&+ \sum_{i=1}^{n-1}\left((\theta-1)\frac{I_{\widehat{p}_i}}{I} + b_{\widehat{p}_i}\right)\left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right) \quad - \gamma(\theta-1)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i} - \gamma\sum_{i=1}^{n-1}\omega_{\widehat{p}_i}\sigma_{c,\hat{p}_i} \\
&+ \sum_{i=1}^{n-1}\left(\frac{1}{2}(\theta-1)(\theta-2)\left(\frac{I_{\widehat{p}_i}}{I}\right)^2 + (\theta-1)\frac{I_{\widehat{p}_i}}{I}b_{\widehat{p}_i} + \frac{1}{2}b_{\widehat{p}_i}^2\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_i} \\
&+ \sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\left(\frac{1}{2}(\theta-1)(\theta-2)\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I} + (\theta-1)\frac{I_{\widehat{p}_i}}{I}b_{\widehat{p}_j} + \frac{1}{2}b_{\widehat{p}_i}b_{\widehat{p}_j}\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j} \\
&+ \sum_{i=1}^{n-1}\frac{1}{2}\left((\theta-1)\left(\frac{I_{\widehat{p}_i\widehat{p}_i}}{I}\right) + b_{\widehat{p}_i\widehat{p}_i} + \right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_i} \quad + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\left((\theta-1)\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + b_{\widehat{p}_i\widehat{p}_j}\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}.
\end{aligned}$$

We approximate the bond price e^{b_t} at each time point t by a multivariate Chebyshev polynomial, i.e.,

$$b_t \;=\; \sum_{j=0}^{n}\alpha_{j,t}^{(T)}T_j(\hat{p}).$$

Note that the PDE for the bond price involves a time derivative. We use an explicit Euler discretization for this time derivative and solve the PDE recursively backwards in time, starting from the boundary condition $b_T = 0$, i.e. $\alpha_{j,T}^{(T)} = 0$ for $j = 0, \ldots, n$.

C.5. Pricing nominal bonds

Let the price of the nominal bond maturing at time T be denoted by $B_t^{T,\$}$ with

$$\begin{aligned} B_t^{T,\$} &= E_t\left[\frac{\xi_T^{\$}}{\xi_t^{\$}}\right] \\ &= E_t\left[\frac{\xi_T}{\xi_t}\frac{e^{\pi_t}}{e^{\pi_T}}\right], \end{aligned}$$

where $\xi_t^{\$} \equiv \xi_t e^{-\pi_t}$ is the nominal pricing kernel. Define $b_t^{\$} \equiv \ln B_t^{T,\$}$ (we again omit the superscript T in the following). Then the Feynman-Kac formula applied to $H(\xi_t^{\$}, b_t^{\$}) = \xi_t^{\$} e^{b_t^{\$}}$ yields the partial differential equation

$$0 \quad = \quad \mathcal{A}H \quad = \mu_{\xi^{\$},t} + \mu_{b^{\$},t} + \frac{1}{2}\frac{d[b_t^{\$}]}{dt} + \frac{d[\xi^{\$}, b^{\$}]_t}{\xi_t^{\$} dt}, \tag{C.6}$$

where $\mu_{\xi^{\$}}$ and $\mu_{b^{\$}}$ are drifts of the processes $\xi_t^{\$}$ and $b_t^{\$}$ respectively. Notice that

$$\frac{d\xi_t^{\$}}{\xi_t^{\$}} = \frac{d\xi_t}{\xi_t} - d\pi_t + \frac{1}{2}d[\pi]_t - \frac{d[\xi, \pi]_t}{\xi_t}.$$

The dynamics of $b_t^{\$}$ are

$$db_t^{\$} = \frac{\partial b_t^{\$}}{\partial t}dt + \sum_{i=1}^{n-1} b_{\widehat{p}_i}^{\$} d\hat{p}_{it} + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1} b_{\widehat{p}_i\widehat{p}_j}^{\$} \sigma_{\hat{p}_i}\sigma_{\hat{p}_j}' dt.$$

Plugging everything into (C.6) yields the following PDE for $b_t^{\$}$:

$$
\begin{aligned}
0 \;=\; & -\beta\theta - (1-\theta)I^{-1} - \gamma \sum_{i=1}^{n} \mu_i^C \widehat{p}_{it} + \frac{1}{2}\gamma^2(\sigma^C)^2 + \frac{\partial b_t^{\$}}{\partial t} - \sum_{i=1}^{n} \mu_i^{\pi} \widehat{p}_{it} + \frac{1}{2}(\sigma^{\pi})^2 + \gamma\rho\sigma^C\sigma^{\pi} \\
& + \sum_{i=1}^{n-1} \left((\theta-1)\frac{I_{\widehat{p}_i}}{I} + b_{\widehat{p}_i^{\$}} \right) \left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1 \\ j\neq i}}^{n} \widehat{p}_{jt}\lambda_{ji} \right) \\
& - \gamma(\theta-1)\sum_{i=1}^{n-1} \frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i} - \gamma \sum_{i=1}^{n-1} b_{\widehat{p}_i}^{\$}\sigma_{c,\hat{p}_i} - (\theta-1)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\sigma_{\pi,\hat{p}_i} - \sum_{i=1}^{n-1} b_{\widehat{p}_i}^{\$}\sigma_{\pi,\hat{p}_i} \\
& + \sum_{i=1}^{n-1} \left(\frac{1}{2}(\theta-1)(\theta-2)\left(\frac{I_{\widehat{p}_i}}{I}\right)^2 + (\theta-1)\frac{I_{\widehat{p}_i}}{I} b_{\widehat{p}_i}^{\$} + \frac{1}{2}(b_{\widehat{p}_i}^{\$})^2 \right) \sigma_{\hat{p}_i}\sigma_{\hat{p}_i}' \\
& + \sum_{i=1}^{n-1}\sum_{j=1}^{n-1} \left(\frac{1}{2}(\theta-1)(\theta-2)\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I} + (\theta-1)\frac{I_{\widehat{p}_i}}{I} b_{\widehat{p}_j}^{\$} + \frac{1}{2} b_{\widehat{p}_i}^{\$} b_{\widehat{p}_j}^{\$} \right) \sigma_{\hat{p}_i}\sigma_{\hat{p}_j}' \\
& + \frac{1}{2}\sum_{i=1}^{n-1} \left((\theta-1)\left(\frac{I_{\widehat{p}_i\widehat{p}_i}}{I}\right) + b_{\widehat{p}_i\widehat{p}_i}^{\$} + \right) \sigma_{\hat{p}_i}\sigma_{\hat{p}_i}' \\
& + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1} \left((\theta-1)\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + b_{\widehat{p}_i\widehat{p}_j}^{\$} \right) \sigma_{\hat{p}_i}\sigma_{\hat{p}_j}'.
\end{aligned}
$$

Again, this PDE involves a time derivative. As for the prices of real bonds, we approximate $e^{b_t^{\$}}$ at each time point t by multivariate Chebyshev polynomials, use an explicit Euler discretization for this time derivative and solve the PDE recursively backwards in time, starting from the boundary condition $b_T^{\$} = 0$.

D. Solution of the Model with Full Information

The dynamics of consumption and the log price level are

$$
\begin{aligned}
d\ln C_t &= \left(\sum_{i=1}^{n} \mu_i^C \mathbb{1}_{\{S_t=i\}} \right) dt + \sigma^C \left(\sqrt{1-\rho^2}\, dW_t^C + \rho\, dW_t^{\pi} \right) \\
d\pi_t &= \sum_{i=1}^{n} \mu_i^{\pi} \mathbb{1}_{\{S_t=i\}} + \sigma^{\pi} dW_t^{\pi},
\end{aligned}
$$

For ease of notation, we relabel the indicator variable as $p_{i,t} = \mathbb{1}_{\{S_t=i\}}$. These are the state variables in a model in which the representative agent perfectly knows the state of the economy at every point in time. The dynamics of these state variables are

$$dp_{i,t} = -p_{i,t} dN_{i,i} + (1 - p_{i,t}) \sum_{j \neq i} dN_{j,i}$$

where the counting process $N_{j,i}$ counts transitions from state j to state i and has the intensity $\lambda_{j,i}$. The counting process $N_{i,i}$ counts the transitions from state i to any other state $j \neq i$ and has the intensity $\sum_{j \neq i} \lambda_{i,j}$. These dynamics imply that the state variables almost surely only take the values 0 and 1.

The indirect utility function of the investor is given by

$$J(C_t, p_{1t}, \dots, p_{nt}) = E_t \left[\int_t^{\infty} f(C_s, J(C_s, p_{1s}, \dots, p_{ns}))\, ds \right].$$

With the same reasoning as in the incomplete information economy (i.e. defining the indirect utility function, setting up the Bellman equation, conjecturing the same functional form for the indirect utility function) and after some simplifications, we arrive at the following algebraic equations (one equation for each state i) for the four unknowns v_i (the log wealth-consumption ratios in each state i):

$$0 \;=\; \left[(1-\gamma)\mu_i^C + \frac{1}{2}(1-\gamma)^2 (\sigma^C)^2 - \beta\theta \right] \;+\; \theta e^{-v_i} + \sum_{j \neq i} \lambda_{ij} \cdot \left(e^{\theta(v_j - v_i)} - 1 \right).$$

The proof is only a very slight modification of the proof in Appendix A of Branger, Kraft, and Meinerding (2016). Similarly, for the log price-dividend ratio ω_i of the equity claim, we get the equations (one for each state i):

$$\begin{aligned} 0 \;=\; & -\beta\theta - (1-\theta)e^{-v_i} + e^{-\omega_i} - \gamma\mu_i^C + \bar{\mu} + \phi\left(\mu_i^C - \bar{\mu}\right) + \frac{1}{2}(\phi - \gamma)^2 (\sigma^C)^2 \\ & + \sum_{j \neq i} \lambda_{ij} \cdot \left(e^{(\theta-1)(v_j - v_i)} \cdot e^{\omega_j - \omega_i} - 1 \right). \end{aligned}$$

For the prices of nominal bonds $b_i^{\$}$ we get the four ordinary differential equations

$$\begin{aligned} 0 \;=\; & -\beta\theta - (1-\theta)e^{-v_i} - \gamma\mu_i^C + \frac{1}{2}\gamma^2 (\sigma^C)^2 + \frac{\partial b_i^{\$}}{\partial t} - \mu_i^{\pi} + \frac{1}{2}(\sigma^{\pi})^2 + \gamma\rho\sigma^C\sigma^{\pi} \\ & + \sum_{j \neq i} \lambda_{ij} \cdot \left(e^{(\theta-1)(v_j - v_i)} \cdot e^{b_j^{\$} - b_i^{\$}} - 1 \right). \end{aligned}$$

Recent Issues

No. 333	Vincent R. Lindner	Solidarity without Conditionality. Comparing the EU Covid-19 Safety Nets SURE, Pandemic Crisis Support, and European Guarantee Fund
No. 332	Gyozo Gyöngyösi, Judit Rariga, Emil Verner	The Anatomy of Consumption in a Household Foreign Currency Debt Crisis
No. 331	Ruggero Jappelli, Loriana Pelizzon, Alberto Plazzi	The Core, the Periphery, and the Disaster: Corporate-Sovereign Nexus in COVID-19 Times
No. 330	Jan Krzyzanowski, Uwe Walz	Bank Regulation, Lending and Patenting: Evidence from the EBA Capital Exercise
No. 329	Jana Eisenkopf, Steffen Juranek, Uwe Walz	Responsible Investment and Stock Market Shocks: Short-Term Insurance and Persistent Outperformance Post-Crisis?
No. 328	Hoang Ha Nguyen Thi, Alfons Weichenrieder	C and S Corporation Banks: Did Trump's Tax Reform Lead to Differential Effects?
No. 327	Fincap team	Non-Standard Errors
No. 326	Wenhui Li, Peter Ockenfels, Christian Wilde	The Effect of Ambiguity on Price Formation and Trading Behavior in Financial Markets
No. 325	Alperen A. Gözlügöl	The Clash of 'E' and 'S' of ESG: Just Transition on the Path to Net Zero and the Implications for Sustainable Corporate Governance and Finance
No. 324	Massimiliano Caporina, Michele Costola	Time-Varying Granger Causality Tests for Applications in Global Crude Oil Markets: A Study on the DCC-MGARCH Hong Test
No. 323	Alexandre Corhay, Thilo Kind, Howard Kung, Gonzalo Morales	Discount Rates, Debt Maturity, and the Fiscal Theory
No. 322	Gianluca Anese, Marco Corazza, Michele Costola, Loriana Pelizzon	Impact of Public News Sentiment on Stock Market Index Return and Volatility
No. 321	Ignazio Angeloni, Johannes Kasinger, Chantawit Tantasith	The Geography of Banks in the United States (1990-2020)
No. 320	Sebastian Steuer, Tobias H. Tröger	The Role of Disclosure in Green Finance