

Nicole Branger - Holger Kraft - Christoph Meinerding

The Dynamics of Crises and the Equity Premium

SAFE Working Paper No. 11

SAFE | Sustainable Architecture for Finance in Europe

A cooperation of the Center for Financial Studies and Goethe University Frankfurt

House of Finance | Goethe University
Grüneburgplatz 1 | 60323 Frankfurt am Main

Tel. +49 69 798 33684 | Fax +49 69 798 33910
info@safe-frankfurt.de | www.safe-frankfurt.de

The Dynamics of Crises and the Equity Premium

Nicole Branger*

Holger Kraft**

Christoph Meinerding**

This version: October 17, 2014

Abstract

There has been a considerable debate about whether disaster models can rationalize the equity premium puzzle. This is because empirically disasters are not single extreme events, but long-lasting periods in which moderate negative consumption growth realizations cluster. Our paper proposes a novel way to explain this stylized fact. By allowing for consumption drops that can spark an economic crisis, we introduce a new economic channel that combines long-run and short-run risk. First, we document that our model can match consumption data of several countries. Second, it generates a large equity risk premium even if consumption drops are of moderate size.

Keywords: General Equilibrium, Asset Pricing, Recursive Preferences, Long-run Risk, Short-run Risk

JEL: G01, G12

*Finance Center Muenster, University of Muenster, Universitaetsstrasse 14-16, 48143 Muenster, Germany. E-mail: nicole.branger@wiwi.uni-muenster.de.

**Goethe University Frankfurt, Faculty of Economics and Business Administration, Grueneburgplatz 1, 60323 Frankfurt am Main, Germany. E-mail: holgerkraft|meinerding@finance.uni-frankfurt.de.

The paper was presented at the 20th Annual Meeting of the German Finance Association in Wuppertal, the 5th Financial Risks International Forum in Paris, the ECB Workshop on “Asset pricing models in the aftermath of the financial crisis” in Frankfurt, the 17th Annual Meeting of the German Finance Association in Hamburg, the 14th Conference of the Swiss Society for Financial Market Research, and at research seminars at the universities of Bielefeld, Frankfurt and Muenster. The authors would like to thank Bjorn Eraker, Stephen Figlewski, Karen Lewis, Claus Munk, Christian Schlag, Ivan Shaliastovich, Mark Trede, Pietro Veronesi, and Amir Yaron as well as conference/seminar participants and discussants for many helpful comments and suggestions. We gratefully acknowledge research and financial support from the Center of Excellence SAFE, funded by the State of Hessen initiative for research LOEWE. Holger Kraft gratefully acknowledges financial support by Deutsche Forschungsgemeinschaft (DFG). This paper was previously circulated under the title “How Does Contagion Affect General Equilibrium Asset Prices?”.

1 Introduction

Modern asset pricing has made significant progress in solving the equity premium puzzle. There is now a variety of approaches that can all generate a sizeable equity risk premium. This is achieved with moderate levels of risk aversion and without implausible implications for the risk-free rate. Two of these approaches are disaster risk models and long-run risk models.¹ Disasters are the most extreme form of short-run consumption risk, which affects the level of consumption, but not its growth rate. Rietz (1988) and Barro (2006) show that even a small probability of large consumption drops can generate a sizeable equity risk premium. Following Bansal and Yaron (2004) the classical example for a long-run-risk variable is a stochastic mean growth rate of consumption. Even small changes in this growth rate constitute a significant risk since they influence the economy over a long time period. Therefore, investors with preferences for early resolution of uncertainty demand a high premium for bearing this risk. Usually, asset pricing models focus either on short-run or long-run risk.² By contrast, the literature has hardly explored the effect of combining these two channels. This is the focus of our paper. We show that a channel which combines short-run and long-run risk adds a new component to the equity risk premium. This component has about the same order of magnitude as the premium for pure short-run risk or pure long-run risk alone. In particular, our approach can generate a realistic equity risk premium even if the possible consumption shocks are of moderate size.

This feature of our model helps to address a fundamental critique on disaster models raised by Constantinides (2008). He challenges the suggestion by Rietz (1988) and Barro (2006, 2009) that the possibility of very large, but rare consumption shocks gives rise to a reasonable risk premium on US equity. Assuming a small unconditional jump probability of 1.7%, but an extreme average jump size of -36% leads to an equity premium of 5.4 percentage points with a risk aversion parameter of 4 in their model. Barro (2006) shows that such disasters can be found in consumption time series of many countries, while the US have been lucky to not experience such an event since World War II. Constantinides (2008) argues that the peak-to-trough method of Barro and Rietz does not match the consumption data. Disasters like World War II do not lead to a single shock of about -40% , but unfold in a series of moderate shocks that add up to an overall drop of -40% . Replacing one large shock by a series of smaller shocks in a model with CRRA preferences

¹An alternative approach uses external habit formation. See Campbell and Cochrane (1999).

²The paper by Wachter (2013) is a notable exception.

reduces the equity premium significantly. As mentioned above, this is not the case in our model, though.

We also address the fundamental critique raised by Julliard and Ghosh (2012). These authors show that, if the rare disaster model of Barro (2006) is used as a data-generating model, the observation of the equity premium puzzle itself would be a rare event. In other words, the observation of the equity premium puzzle is due to a so-called *Peso problem*³, which means that the US consumption path over the past 60 years has just been a very unlikely draw from a distribution with fat tails. We thus simulate our model and compute the mean excess return of the equity claim including all negative shocks in the sample. We find that the new channel of our model allows us to perfectly match the mean excess return. This remains true even in samples where we observe several downward jumps in stock prices. Our model can thus address the equity premium puzzle without relying on a *Peso problem* explanation.

To describe our model in more detail, notice that it involves three types of risk channels: short-run risk, long-run risk, and a combination of both. The short-run risk channel is modeled via consumption drops that can happen with a certain probability. This channel affects the level, but not the growth rate of consumption. Its impact is instantaneous. The probability of a consumption drop is assumed to depend on the state of the economy (good, bad) modeled by a two-state Markov chain. It is larger in the bad state. This constitutes the long-run risk channel that affects the growth rate, but not the level of consumption. Its effect is long-lasting. As a third and novel channel, we propose a blend of both channels that has short-run and long-run implications. We model this idea by introducing jump-induced regime switches that combine a consumption drop with a transition from the good into the bad state. In other words, there are certain consumption drops in normal times (good state) that spark an economic crisis (bad state). As a consequence, short-run risk earns an additional long-run-risk premium, and long-run risk earns an additional short-run risk premium. These effects do not simply add up, but reinforce each other.

Following Duffie and Epstein (1992a), Bansal and Yaron (2004) and Wachter (2013), among others, we study an economy that is populated by a representative investor with recursive utility. His risk aversion and intertemporal elasticity of substitution are both larger than one. This implies that the investor has a preference for early resolution of uncertainty and demands a risk premium for both short-run and long-run risk. We solve for the equilibrium quantities in closed form, in particular for the equity risk premium.

³See, e.g., Goetzmann and Jorion (1999).

This allows us to decompose the equity risk premium into its three components (short run, long run, combined). To document that our new channel is of first-order importance for the sizes of the equity risk premium and risk-free rate, we compare our model with *jump-induced regime switches* to the following well-established alternatives: First, we consider a model with *separated regime switches* where short-run risk is decoupled from long-run risk in the sense that transitions to the bad state and consumption drops happen separately. Second, we study a *peak-to-trough* model that involves only short-run risk. Therefore, the whole dynamics of an economic crisis are condensed to one consumption drop which is assumed to happen at a single point in time.⁴

We calibrate our model to consumption growth rates of 42 different countries.⁵ Our focus is on matching both the unconditional distribution of annual consumption declines and the clustering of downward jumps in bad times. This gives rise to a moderate jump size of -3% to -7% and an unconditional jump intensity between 0.14 and 0.42. Our model explains the left tail of the unconditional distribution of annual consumption growth rates almost perfectly. It also reproduces the distributions of the lengths and sizes of crises in the data. The hypothesis that the annual consumption data has been generated by our model cannot be rejected. For a calibration with a jump size of -5% , the equity risk premium in the good state is 6.94 percentage points. The premium for jump-induced regime switches from the good to the bad state is 4.95 percentage points, which comprises a premium of 0.28 percentage points for short-run jump risk, 2.22 percentage points for long-run risk, and 2.45 percentage points for the combination of long- and short-run risk. This last term is solely present in our model. In the model with separated regime switches, the equity risk premium in the good state is only 4.99 percentage points. In the peak-to-trough model with the same local consumption distribution, the premium is 2.9 percentage points. Notice that one can boost the values of the equity premium in the model with separate regime switches and the peak-to-trough model. By allowing for extreme consumption drops (peak-to-trough calibration), one can reach a higher equity premium. But then the Constantinides (2008) critique applies and the calibration of the consumption data is off. To summarize, only our model is able to achieve a realistic fit of the consumption data and at the same time generate a realistic equity premium.

Finally, several robustness checks document that our results are not driven by one partic-

⁴Notice that all three models are silent on the economic reasons for crises. In this sense, they are reduced-form approaches. For clarity, we also abstract from further risk factors like a stochastic drift or stochastic volatility, which would distract from our new channel and obscure our main point.

⁵We thank Robert Barro for providing this dataset on his webpage.

ular calibration of the model. The consumption data can be matched using different jump sizes between -3% and -7% . All calibrations lead to a sizeable risk premium for our new channel combining long-run and short-run risk. Regarding the preference parameters, we show that the elasticity of intertemporal substitution (EIS) plays a key role in our model. For low levels of the EIS where the investor no longer has a pronounced preference for early resolution of uncertainty, the extra risk premium for jump-induced regime switches shrinks dramatically and can even become negative. In particular, in the special case of constant relative risk aversion (CRRA), the total risk premium is higher in a model with separated regime switches than in a model with jump-induced regime switches. These findings are in line with the long-run-risk literature (see, e.g., Bansal and Yaron (2004)).

Our paper is related to the literature on asset pricing with rare disasters. Rietz (1988) and Barro (2006, 2009) entertain the disaster risk explanation to rationalize a high equity risk premium. Extensions of the basic model have been studied by Chen, Joslin, and Tran (2012) and Julliard and Ghosh (2012), among others. Gabaix (2012) analyzes a model with time-varying jump intensities. As a response to the critique by Constantinides (2008), Barro, Nakamura, Steinsson, and Ursua (2013) propose and estimate a complex macroeconomic model with recursive preferences where disasters unfold over several years including a recovery period afterwards. Their paper is similar in spirit to ours, but the mechanics of our model are simpler. For instance, this allows us to disentangle the implications of short-run and long-run risk in closed form and to explicitly assess the impact of our combined channel of long-run and short-run risk. Similar to the critique of Constantinides (2008), the assumption of extreme jumps is also questioned by Backus, Chernov, and Martin (2011). They use index options to estimate the implied distribution of consumption jumps and document that option prices imply more frequent but less extreme outcomes than needed in disaster risk explanations of the equity premium. Finally, several papers study the implications of jump risk and/or regime switching processes in long-run risk models. Hung (1994) studies a model with Epstein-Zin utility in which market fundamentals follow a bivariate Markov switching process. His setting is similar to the model with separated regime switches that we consider as one possible benchmark. Benzoni, Collin-Dufresne, and Goldstein (2011) analyze a model with jumps in the expected growth rate and volatility and show that this model provides an equilibrium foundation of the volatility smile. Drechsler and Yaron (2011) analyze the variance risk premium in a long-run risk model with time-varying jump risk. Wachter (2013) proposes a time-varying probability of rare disasters which can explain the high stock market volatility. Tsai and Wachter (2014) apply this model to growth and value stocks.

The remainder of the paper is organized as follows. Section 2 describes the different model setups used in this paper. Section 3 provides details on the calibration and a comparison with the stylized facts about consumption drops. Section 4 determines the equilibria of the different models. Additionally, we study the equity risk premium and its components in regime switching as well as disaster models. Section 5 performs robustness checks. Section 6 concludes. All proofs are in the Appendix.

2 Consumption Dynamics

2.1 Model with Jump-induced Regime Switches

We consider a continuous-time Lucas tree economy with an infinite horizon. There is one tree generating a perishable consumption good which serves as numeraire. The economy can be in either of two states which we denote by g and b (for ‘good’ and ‘bad’). The states are formally captured by a Markov chain Z . Conditional on the state of the economy $Z_t \in \{g, b\}$, the outcome of the tree follows a jump-diffusion process. If the economy is in the good state, then

$$\frac{dC_t}{C_{t-}} = \mu dt + \sigma dW_t + LdN_t^{g,g} + LdN_t^{g,b}. \quad (1)$$

In the bad state, the dynamics are

$$\frac{dC_t}{C_{t-}} = \mu dt + \sigma dW_t + LdN_t^{b,b}, \quad (2)$$

where W is a Wiener process. The jump processes $N^{g,g}$ and $N^{b,b}$ count downward consumption jumps in the good and the bad state that do not have an effect on the regime. We refer to these drops as ‘normal jumps’. A transition from the good to the bad state is captured by the jump process $N^{g,b}$. The decisive feature of our model is that this change comes together with a downward jump in consumption, i.e. we allow for the joint occurrence of negative consumption realizations (‘cash flow shocks’) and changes in the distribution of future consumption growth (‘regime shifts’). Regime switches from the bad to the good state are triggered by a fourth counting process $N^{b,g}$. These jumps only change the state, but do not directly affect the level of consumption. For ease of exposition, we do not include changes from the good to the bad state which are not linked to consumption jumps.

For simplicity, the drift μ , volatility σ and jump size $L < 0$ are constant and state-independent.⁶ The jump intensity depends on the state Z of the economy. The parameter $\lambda^{g,b}$ captures the probability of a jump-induced regime switch in normal times, whereas $\lambda^{g,g}$ and $\lambda^{b,b}$ denote the intensities for jumps that do not trigger regime switches. In line with the interpretation of the states as 'good' and 'bad', we assume that consumption shocks are more frequent in the bad state than in the good state, i.e. $\lambda^{g,g} + \lambda^{g,b} < \lambda^{b,b}$.

2.2 Model with Separated Regime Switches

To tease out the effect of jump-induced regime switches, we also consider a model with regime switches that do not induce immediate negative consumption realizations. We refer to these regime changes as 'separated regime switches'. The corresponding consumption dynamics become

$$\frac{dC_t}{C_{t-}} = \mu dt + \sigma dW_t + L dN_t^{g,g} \quad (3)$$

in the good state and

$$\frac{dC_t}{C_{t-}} = \mu dt + \sigma dW_t + L dN_t^{b,b}, \quad (4)$$

in the bad state, respectively. The counting process $N^{g,b}$ still triggers a regime change from the good to the bad state, but does not lead to an immediate consumption shock. When we compare jump-induced and separated regime switches, we assume that the local distribution of consumption and the transition probabilities of the Markov chain are the same in both models. The jump intensities for the regime switches $N^{g,b}$ and $N^{b,g}$ are thus identical, while the intensity of $N^{g,g}$ is larger with separated regimes than with jump-induced regimes.

2.3 Peak-to-Trough Model

The second model that we use for comparison abstracts from different states and assumes that consumption growth is i.i.d. The consumption dynamics are given by

$$\frac{dC_t}{C_{t-}} = \mu dt + \sigma dW_t + L dN_t, \quad (5)$$

where N is a Poisson process with intensity λ . This model represents a special case of the Barro (2006) model with constant jump size. It can explain the equity risk premium if

⁶These assumptions can be relaxed, but this would not add additional insights concerning our main focus.

jumps are disastrous but rare events. Barro (2006) calibrates the jump parameters to a panel of consumption data from different countries and argues that there is indeed evidence for disasters. An economic crisis is assumed to be a single (very) negative consumption realization. The size of this consumption shock is usually calibrated to the observed consumption drops from peak to trough. Crisis which may last over several years are thus mapped into a single event. The dynamics of crises are disregarded, which is criticized by Constantinides (2008).

We consider two calibrations of this model. In the first calibration ('local calibration'), we stick to the unconditional local distribution of consumption growth rates from the other two models. Consequently, we set the jump intensity λ equal to the unconditional average jump intensity and keep the jump size L fixed at the same value as in the other models. The second calibration ('disaster calibration') follows the idea of Rietz (1988) and Barro (2006) who model disasters as rare but extreme events. We set the jump size L equal to the average consumption drop during a crisis from peak to trough. The jump intensity in the disaster calibration is set equal to the probability of observing an extreme crisis in the model with jump-induced regime switches. We thus condense a crisis which usually stretches over several years to a single event. Consequently, under the disaster calibration, jumps are less frequent but more severe than under the local calibration.

3 Data and Calibration

This section describes the data and explains how we calibrate the different models. First we calibrate our model with jump-induced regime switches so that it matches both the unconditional distribution of annual consumption growth rates and the durations of crises. Naturally, we focus on the left-hand side of the distribution which contains information about the nature of consumption crises. Then we turn to the model with separated regime switches. This model is identical to our model except that jump-induced regime switches are not possible. To tease out the effect of this channel, we take the calibration of the jump-induced regime switching model as given, switch off the jump-induced regime switches and then adjust the calibration such that the local distribution of consumption growth rates and the transition probabilities are the same as in the model with jump-induced regime switches.⁷ Finally, we consider the peak-to-trough model and provide two calibrations. In

⁷Alternative calibrations could either fit the consumption data or the equity premium, but not both. See also Section 3.3.

the first one, we match the local distribution of consumption, whereas in the second one we apply the peak-to-trough approach by Barro (2006).

3.1 Data

All calibrations are based on the consumption dataset that is available on Robert Barro's webpage. This dataset contains annual consumption growth rates of 42 different countries ranging (at most) from 1791 to 2009. Altogether, the dataset consists of 4,933 country-year consumption growth observations. The dataset also contains GDP growth rates for all the countries. Neither the parameters nor the numerical results will change significantly if we use GDP growth rates. Therefore, we do not report the corresponding results here. To mitigate potential Peso problems as discussed in Goetzmann and Jorion (1999), we use the full dataset rather than just US consumption data. For the same reason, we also do not exclude data points from the sample. We pool the consumption data from all countries into one single time series by adding country after country.⁸ We exclusively use consumption data to calibrate the model and do not rely on asset price data in this section.

The row labeled 'Data' in Panel A of Table 1 reports the unconditional moments. The mean annual consumption growth rate is 0.02 and thus similar to the value observed for U.S. consumption data to which asset pricing models are usually calibrated. The standard deviation is 0.065 and significantly exceeds the value for U.S. data of around 0.02. The distribution is slightly right-skewed with a skewness of 0.37, and highly leptokurtic with a kurtosis of 10.04. Since the left tail of the growth rate distribution contains information about consumption crises, Panel B reports values of the empirical unconditional cumulative distribution function below zero. For instance, the probability that the annual consumption growth is below 0 (−10%) amounts to around 30% (3.3%).

Constantinides (2008) points out that understanding the part of the equity risk premium which can be attributed to potential consumption disasters requires a dynamic analysis. In particular, the duration of crises is of crucial importance, since annual consumption growth over the peak-to-trough period of a disaster is highly autocorrelated. Therefore, Table 2 provides some evidence on the durations of crises. We define the length of a crisis as the number of consecutive years where the consumption growth rates are below a given threshold. For a threshold of 0%, there are 1,492 observations with consumption growth

⁸In our calibrations, we make sure that we do not create 'artificial' crises across countries through this.

rates below 0%. This gives rise to 906 crises periods of different lengths. Panel A of Table 2 reports the unconditional probabilities that a crisis in the sample has a length of one up to ten years. For instance, 61.6% of all crises in the data last only one year, 23% have a duration of two years and so on. Finally, there is exactly one recession of 10 years in the dataset (India in the 1930s). Panels B and C report the corresponding probabilities for more pronounced annual consumption drops of at least -5% and -10% , respectively. Even if consumption growth were i.i.d., crises can last for several years. Therefore, clustering per se should not be interpreted as evidence that observed crises are generated by a model with regimes. To study whether consumption jumps are indeed more frequent than in models with i.i.d. consumption growth, we calculate various conditional probabilities that bad years follow an initial bad year. We then check whether these conditional probabilities are larger than the corresponding unconditional probabilities. We formalize this idea by studying sequences of Bernoulli trials. The year- t outcome X_t is 0 if “consumption growth is above the threshold” and 1 if “consumption growth is below the threshold”. We consider thresholds of 0%, -5% , or -10% , respectively. The numbers in Table 2 are the probabilities for a certain duration of a crisis given that year t has been the beginning of the crisis. For a threshold of 0%, the first four conditional probabilities are given by (see first line in Panel A of Table 2):

$$\begin{aligned}
\text{Prob}(X_{t+1} = 0|X_t = 1, X_{t-1} = 0) &= 0.616 & (6) \\
\text{Prob}(X_{t+2} = 0, X_{t+1} = 1|X_t = 1, X_{t-1} = 0) &= 0.230 \\
\text{Prob}(X_{t+3} = 0, X_{t+2} = 1, X_{t+1} = 1|X_t = 1, X_{t-1} = 0) &= 0.092 \\
\text{Prob}(X_{t+4} = 0, X_{t+3} = 1, X_{t+2} = 1, X_{t+1} = 1|X_t = 1, X_{t-1} = 0) &= 0.038
\end{aligned}$$

For a threshold of 0%, the empirical unconditional probability of being in a crisis is $\text{Prob}(X_t = 1) = 0.302$. Therefore, if the data were generated by a sequence of *independent identical* Bernoulli trials⁹, then the probabilities (6) could be calculated in the following way:

$$\begin{aligned}
\text{Prob}(X_t = 0) &= 0.698 \\
\text{Prob}(X_t = 0)\text{Prob}(X_t = 1) &= 0.698 \cdot 0.302 = 0.211 \\
\text{Prob}(X_t = 0)\text{Prob}(X_t = 1)^2 &= 0.698 \cdot 0.302^2 = 0.064 \\
\text{Prob}(X_t = 0)\text{Prob}(X_t = 1)^3 &= 0.698 \cdot 0.302^3 = 0.019
\end{aligned}$$

i.e. only the unconditional probabilities of (not) being in a crisis would matter. Compared to a model with i.i.d. consumption realizations, crises of length 3 years are 50% more likely

⁹Notice that this implies $\text{Prob}(X_{t+k} = n) = \text{Prob}(X_t = n)$ for all $k = 1, 2, \dots$ and $n \in \{0, 1\}$.

in the data and crises of length 4 years are twice as likely. These results are in line with Constantinides (2008) and provide evidence that consumption drops cluster in reality.

3.2 Model with Jump-induced Regime Switches

We calibrate our model so that it matches both the unconditional distribution of annual consumption growth rates and the conditional distribution of the durations of crises. The consumption dynamics in our model depend on seven parameters (μ , σ , L , and four intensities $\lambda^{j,k}$). As we are going to show in the following, there is a range of plausible parameter combinations that can match the data well. In particular, the fit is robust with respect to choosing the loss size L . We calibrate the model in several steps: First, we fix $\sigma = 0.04$. This choice makes sure that the probability for negative consumption growth rates is already of the right order without jumps. Additionally, jumps are essential to generate the right amount of severe negative outcomes. The loss size L is however hard to estimate. Therefore, we study five alternative calibrations each involving a different value of L . More precisely, we consider loss sizes of $L = -0.03, -0.04, \dots, -0.07$ and choose the jump intensities $\lambda^{j,k}$ accordingly. Finally, we fix the drift parameter μ so that the corresponding calibration matches the unconditional expected consumption growth.

Notice that there are trade-offs between the jump size and the jump intensities that can be used to calibrate the jump intensities. A given *unconditional* distribution of annual consumption growth rates can be fitted by a small value of L (e.g. -0.03) and a high value of the average jump probability, i.e. large values of $\lambda^{g,g} + \lambda^{g,b}$ and $\lambda^{b,b}$, or a large value of L and a small average jump probability. The empirical properties of crises then put restrictions on the relation between $\lambda^{b,b}$ and $\lambda^{b,g}$. If the *conditional* jump probability in the bad state is low (small $\lambda^{b,b}$), we need to extend the average duration of the bad regime, i.e. reduce $\lambda^{b,g}$. Otherwise, we would not have sufficient autocorrelation in our model-generated data. The unconditional probability to be in a bad state puts an additional restriction on the relation between $\lambda^{g,b}$ and $\lambda^{b,g}$. A small value of $\lambda^{b,g}$ implies a small value of $\lambda^{g,b}$ and vice versa. Finally, for given values of $\lambda^{g,b}$ and L , we choose the intensity $\lambda^{g,g}$ so that the sum $\lambda^{g,g} + \lambda^{g,b}$ is in the right range.

The columns labeled ‘Jump-induced regime switches’ of Table 3 report the five different calibrations described above. The drift rate μ is about 0.03. The jump intensity $\lambda^{g,g}$ is 0.02 in all specifications. The other intensities decrease if L gets more extreme: $\lambda^{g,b}$ ranges from 0.12 to 0.04, $\lambda^{b,b}$ from 2.9 to 0.9, and $\lambda^{b,g}$ from 1.08 to 0.36. In all specifications, the values of $\lambda^{b,g}$ and $\lambda^{g,b}$ imply unconditional probabilities of the two states of 90% and

10%, respectively. Notice that the probability of entering a crisis becomes smaller if the loss size L is larger.

To document that our calibrations can explain the consumption data, we focus on the above-mentioned dimensions of the data: the unconditional distribution of consumption growth rates and the dynamics of crises. For each of the five parameterizations in Table 3, we simulate 500,000 years of daily consumption realizations and aggregate the data to annual time series of consumption growth rates. Table 1 reports the unconditional moments and tail probabilities of these simulated annual consumption growth rates. In all specifications, the unconditional annual mean consumption growth rate is 0.02, exactly as in the data. The unconditional annual standard deviation of consumption is between 0.051 and 0.055 for all specifications. This is slightly below the value of 0.065 in the data. Furthermore, Panel B documents that the model fits the left tail almost perfectly.

To formally assess the fit of the simulated to the empirical data, we perform a Kolmogorov-Smirnov (KS) test. Since our focus is on the left tail of the distribution, we truncate the distribution at a level of 0% (and also -2% and -4% as a robustness check). The null hypothesis of the two-sample Kolmogorov-Smirnov test is that two data samples have been generated by the same (unconditional) distribution. In our data, the p -value from the KS test is always above 0.02 (and sometimes as high as 0.48 for a truncation level of -4% and $L = -0.05$). Therefore, we cannot reject the null hypothesis at the 1% level that a model with jump-induced regime switches has generated the empirical consumption data.¹⁰

Our second objective is to match the dynamics of crises described by their lengths. For all five parametrizations, Table 2 reports the distributions of the lengths of crises. In all cases, our model matches the data well. Notice that for a threshold of -10% there are 132 crises only. All except of one have a duration of less than 4 years. The empirical probability of 0.008 that a crisis lasts for exactly 4 years is thus due to a single crisis (Netherlands 1940-1943).

To summarize, for all five parameterizations the left tail of the simulated consumption distribution is not statistically significantly different from the tail of the empirical distribution. Besides, the average durations of crises are similar to the durations in the empirical data. Notice that we focus on crises, i.e. on the left tail of the consumption distribution.

¹⁰As a robustness check, we also performed KS tests at the country level, i.e. we compared the model-generated time series to each of the 42 country-specific time series of consumption growth rates in the dataset. For a truncation level of 0%, the p -values are above 0.01 for 37 countries and above 0.1 for 25 countries.

By including features such as stochastic volatility, stochastic growth rates, or stochastic jump sizes, one could also try to fit the right tail of the distribution. This is beyond the scope of this paper.

3.3 Model with Separated Regime Switches

The model with separated regime switches (3) and (4) does not allow for jump-induced regime changes. We choose the parameters such that the local distribution of consumption growth rates and the transition probabilities are the same as in the model with jump-induced regime switches. The columns labeled “Separated regime switches” of Table 3 summarize the five resulting calibrations. The parameters μ , σ and L remain unchanged. Also the probability of consumption jumps in the good state is the same as before. Therefore, the value of $\lambda^{g,g}$ is now equal to the sum of $\lambda^{g,g}$ and $\lambda^{g,b}$ in the model with jump-induced regime switches (e.g., we set $\lambda^{g,g} = 0.02 + 0.12 = 0.14$ for $L = -0.03$ in Case I). The other jump intensities remain unchanged.

Crises periods no longer start with an initial drop in consumption, which implies that the total consumption loss over a crisis is smaller. This also materializes in the simulations of the five model calibrations. The probabilities for extreme crises are smaller, and the measured durations of crises are also smaller than before. In line with these findings, the KS test now rejects the hypothesis that the models with separated regime switches have generated the empirical consumption data.

Alternatively, we could also calibrate the model such that it fits the consumption data as close as possible. Results not reported here, show that one can find calibrations that the KS test cannot reject. These calibrations differ only very slightly from the ones presented above. One would have to set the volatility parameter σ to 0.042 instead of 0.04, increase the jump intensities by a small amount etc. However, such calibrations generate an equity premium that is still too low (about 4–5 percentage points). Only our model with jump-induced regime switches is able to achieve a realistic fit of the consumption data and at the same time generate a realistic equity premium.

3.4 Peak-to-Trough Model

For the peak-to-trough model (5), we consider two calibrations. In the first calibration (reported in the columns labeled ‘local calibration’ of Table 3), we match the local distribution of consumption. The jump intensity λ is set to the unconditional average jump

intensity in the model with jump-induced regime switches. The jump size L and the volatility σ have the same values as in the previous models. Not surprisingly, simulations show that the model fails to match the distribution of consumption growth rates over one year and over several years. In particular, the model fails to generate enough long-lasting crises of two or more consecutive years with consumption growth below -5% or -10% . Nevertheless, we include this model in our asset pricing analysis. This allows us to not only study the impact of jump-induced regime switches on asset prices, but also the effect of regime switches per se. The local calibration of the peak-to-trough model serves as one benchmark for this analysis.

The second calibration follows the peak-to-trough calibration approach of Barro (2006). We use the model with jump-induced regime switches and simulate 500,000 years of consumption realizations for each of our five parametrizations. We then extract all crises with one or several consecutive years of negative consumption growth. The peak-to-trough consumption loss is the total drop in consumption over these years. The number of crisis periods in our simulated samples is roughly 98,000 for all five parametrizations. Figure 1 depicts the resulting histogram of peak-to-trough disaster sizes for Case III. We then extract all crises for which the peak-to-trough disaster size exceeds -15% , which is also the threshold used by Barro (2006). This leaves us with about 9,000 – 12,000 disasters, depending on the parametrization. For the peak-to-trough ‘disaster calibration’, we set the jump intensity λ equal to the number of disasters divided by 500,000 and the jump size equal to the average peak-to-trough consumption loss. The column ‘disaster calibration’ of Table 3 reports the exact parameter values.

The simulated time series generated from this disaster calibration are similar to the data-generating processes used by Barro (2006) except that we assume a constant jump size. It is not surprising that the disaster calibration of the peak-to-trough model is not able to match the stylized facts of the consumption data. For instance, the KS tests reject the peak-to-trough model for all values of L . Similar as the model with separated regime switches, the peak-to-trough model cannot resolve the tradeoff between matching consumption data and generating a sizeable equity premium.

4 Asset Pricing

We now turn to the asset-pricing implications of our model. In this section, we focus on the parametrization with a loss size of $L = -0.05$ (Case III). Section 5 analyzes the

alternative parametrizations.

4.1 Representative Investor

Our economy is populated by a representative investor with stochastic differential utility as introduced by Duffie and Epstein (1992b). His subjective time preference rate is β , his relative risk aversion is γ , and his elasticity of intertemporal substitution is ψ . The investor has an infinite planning horizon, and his indirect utility function is

$$J_t = E_t \left[\int_t^\infty f(C_s, J_s) ds \right],$$

where the aggregator f is given by

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

and $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$. Following Bansal and Yaron (2004), among others, we assume $\gamma > 1$ and $\psi > 1$. Therefore, the investor has a preference for early resolution of uncertainty. In the numerical examples, we first assume $\beta = 0.03$, $\psi = 2$, and $\gamma = 6$. As a robustness check, we also study alternative values of the preference parameters including the CRRA case, $\psi = \frac{1}{\gamma}$, where the investor is indifferent towards the resolution of uncertainty.

4.2 Pricing Kernel

Following Duffie and Epstein (1992a,b) and Benzoni, Collin-Dufresne, and Goldstein (2011), we solve for the pricing kernel.¹¹ The agent's indirect utility J is

$$J_t = \frac{C_t^{1-\gamma}}{1-\gamma} \beta^\theta e^{\theta v^{Z_t}}.$$

The pricing kernel ξ is given by

$$\xi_t = \beta^\theta C_t^{-\gamma} e^{-\beta\theta t + (\theta-1) \left(\int_0^t e^{-v^{Z_u}} du + v^{Z_t} \right)}, \quad (7)$$

where v^{Z_t} is the logarithm of the wealth-consumption ratio.¹² It depends on the state of the economy Z and can thus take only two values, v^g and v^b , which solve the following

¹¹Details of the derivation as well as the proofs of all following results can be found in Appendix A.

¹²See, e.g., Campbell, Chacko, Rodriguez, and Viceira (2004) and Benzoni, Collin-Dufresne, and Goldstein (2011).

system of equations:

$$\begin{aligned}
0 &= e^{-v^g} - \beta + \left(1 - \frac{1}{\psi}\right) \mu - \frac{1}{2} \gamma \left(1 - \frac{1}{\psi}\right) \sigma^2 \\
&\quad + \frac{1}{\theta} \lambda^{g,g} [(1+L)^{1-\gamma} - 1] + \frac{1}{\theta} \lambda^{g,b} [(1+L)^{1-\gamma} e^{\theta(v^b-v^g)} - 1] \\
0 &= e^{-v^b} - \beta + \left(1 - \frac{1}{\psi}\right) \mu - \frac{1}{2} \gamma \left(1 - \frac{1}{\psi}\right) \sigma^2 \\
&\quad + \frac{1}{\theta} \lambda^{b,b} [(1+L)^{1-\gamma} - 1] + \frac{1}{\theta} \lambda^{b,g} [e^{\theta(v^g-v^b)} - 1].
\end{aligned}$$

The wealth-consumption ratios are $e^{v^g} = 32.65$ in the good state and $e^{v^b} = 30.25$ in the bad state. In line with intuition, the agent's total wealth is smaller in the bad state, i.e. $v^g > v^b$.

4.3 Price-Dividend Ratios

Given the pricing kernel (7), we can price all claims in the economy including the dividend claim. Following Bansal and Yaron (2004) and Wachter (2013), among others, dividends D_t are modeled as a levered claim to consumption. The dividend dynamics are

$$\frac{dD_t}{D_{t-}} = \mu dt + \phi \sigma dW_t + [(1+L)^\phi - 1] dN_t^{g,g} + [(1+L)^\phi - 1] dN_t^{g,b}$$

in the good state and

$$\frac{dD_t}{D_{t-}} = \mu dt + \phi \sigma dW_t + [(1+L)^\phi - 1] dN_t^{b,b}$$

in the bad state. We assume a leverage parameter $\phi = 2$ for both diffusion risk and jump risk.¹³ Following Longstaff and Piazzesi (2004), we choose the same drift rate μ for consumption and dividends. The pricing equation for the dividend claim is

$$D_t e^{w_t} = \mathbb{E}_t \left[\int_t^\infty \frac{\xi_\tau}{\xi_t} D_\tau d\tau \right],$$

where w is the logarithm of the price-dividend ratio. Analogously to the log wealth-consumption ratio, the log price-dividend ratio w depends on the state of the economy. Its two possible values, w^g and w^b , satisfy a system of equations that is provided in Appendix A.

¹³In the literature, one can find different choices of the leverage parameter. Barro, Nakamura, Steinsson, and Ursua (2013) assume a rather low leverage of $\phi = 1.5$, whereas Bansal and Yaron (2004) set $\phi = 3$. In the robustness checks, we will study alternative parametrizations.

In our model with jump-induced regime switches, the price-dividend ratios are $e^{w^g} = 11.91$ in the good state and $e^{w^b} = 9.38$ in the bad state. Upon a jump from the good to the bad state, the price-dividend ratio drops by 21.16%. In the model with separated regime switches, the price-dividend ratios are $e^{w^g} = 13.24$ and $e^{w^b} = 10.13$, and the drop in the price-dividend ratio is 23.49%. Compared to our model the higher price-dividend ratios reflect the lower overall risk in this economy. Regime switches from the good to the bad state are disentangled from consumption losses. Since the agent is less averse to two small downward jumps in prices than to one large jump, prices are higher. With less risk in the good state, the price difference has to be larger in the model with separated regime switches, too.

In the peak-to-trough model, the price-dividend ratio is constant. It is equal to 19.31 for the local calibration and 17.19 for the disaster calibration. The price-dividend ratios are larger than in each of the models with regimes. Since potential consumption jumps are more severe, the peak-to-trough disaster calibration exhibits more short-run risk than the previous models. There is however no long-run risk. Additional short-run risk lowers the price-dividend ratio, while omitting long-run risk has the opposite effect. Our results show that the impact of long-run risk on the price-dividend ratio dominates the impact of short-run risk.

4.4 Equity Risk Premium

Table 4 reports the local equity risk premia resulting from the different approaches. In the model with jump-induced regime switches, the local equity risk premium is 0.0694 in the good state and 0.1956 in the bad state. Since the economy is in the good state 90% of the time, this yields an unconditional equity risk premium of 0.0820. The model with separated regime switches generates an equity risk premium of 0.0499 in the good state and 0.2164 in the bad state, leading to an unconditional average of 0.0665. Finally, the peak-to-trough model implies an equity risk premium of 0.0290 for the local calibration and 0.0671 for the disaster calibration.

The local calibration of the peak-to-trough model thus results in the lowest equity risk premium. The value of 0.029 is significantly smaller than the typical values of 0.05 – 0.08 estimated from the data. The disaster calibration combines several small jumps to one large jump, i.e. it amplifies the short-run risk. Barro (2006) shows that this crisis feature increases the equity risk premium heavily. By contrast, the model with separated regime switches involves long-run risk. As Bansal and Yaron (2004) document, changes in future

consumption dynamics constitute a significant risk, for which an agent with a preference for early resolution of uncertainty demands a high equity risk premium. Our model with jump-induced regime switches entertains a third channel combining long- and short-run risk. The extra risk premium arising from this new channel has about the same order of magnitude as the premium for pure long-run risk.

Since we can solve our model in closed form, we can explicitly decompose the equity risk premium into its components. This allows us to explicate the mechanism through which our model generates a large equity premium. The instantaneous asset return is defined by $dR = dP/P + D/P dt$. The expected excess return on the dividend claim can be decomposed into diffusion and jump risk premia:

$$\frac{E_t[dR_t^i]}{dt} - r_f^i = RP^{diff} + \sum_{k=g,b} RP^{i,k}.$$

where $i \in \{g, b\}$ denotes the state of the economy. The diffusion risk premium is given by

$$RP^{diff} = \phi\sigma \cdot \gamma\sigma.$$

The first term captures the exposure of the price to diffusive risk. This coincides with the exposure of the dividend to diffusive risk, since there is no diffusive state variable in the models under consideration. The second term is the market price of risk which follows from the dynamics of the pricing kernel. As the diffusion parameters are the same across all specifications, the diffusion risk premium is equal to 0.0192 in all models.

The differences in the unconditional equity risk premia across models and parametrizations arise from the jump risk premia. Each jump risk factor carries a jump risk premium, i.e. there is one jump risk premium in the peak-to-trough model and two jump risk premia in every state of the models with regimes. Table 4 reports the components of the conditional risk premia in all models. Since the good state prevails 90% of the time, the unconditional risk premia mainly depend on the risk premia in the good state. Taking this into account our numerical results show that the premium for jump-induced regime switches is quantitatively the most important component of the unconditional risk premium. This holds true for our model, but also compared to the components of all other model specifications.

To shed more light on this finding, we explicitly decompose the jump risk premia further into their constituents. In general, the risk premium for a jump risk factor with constant jump size and intensity is the product of three figures:

$$RP = \lambda \cdot \eta \cdot \zeta.$$

The amount of jump risk depends on the intensity λ and the exposure ζ . This exposure is defined as the relative change of the asset price upon the corresponding event. It follows from the equilibrium price dynamics. The market price of risk η measures the equilibrium compensation for one unit of this risk. It follows from the dynamics of the pricing kernel. Table 5 reports the intensities, market prices of risk and exposures for all jump risk factors and all model specifications.¹⁴

In general, the market prices of risk and the exposures can involve a short-run and a long-run component. An event which changes the consumption by X and/or switches the state (and thus affects the valuation ratios) gives rise to the following market price of risk:

$$\eta = 1 - \underbrace{(1 + X)^{-\gamma}}_{\text{SRR}} \cdot \underbrace{e^{(\theta-1)(v \text{ after jump} - v \text{ before jump})}}_{\text{LRR}}.$$

The immediate consumption drop upon the event gives rise to a short-run component $(1 + X)^{-\gamma}$. If $X < 0$, the market price of jump risk $1 - (1 + X)^{-\gamma}$ for jumps that do not change the regime ('normal jumps') is negative. The long-run component captures the compensation for regime switches. A change in the future consumption dynamics affects the agent's continuation utility and thus the equilibrium wealth-consumption ratio. Since $\theta < 1$, the long-run component makes the market price of risk more negative if the wealth-consumption decreases upon the jump. The change in the wealth-consumption ratio is priced if the agent is not indifferent towards the timing of the resolution of uncertainty ($\theta \neq 1$). In the special case of CRRA preferences, the market price of jump risk involves a compensation for short-run consumption risk only.

Similarly, the exposure of the dividend claim to an event that changes the consumption by X and/or switches the state can be decomposed as follows:

$$\zeta = \underbrace{(1 + X)^\phi}_{\text{SRR}} \cdot \underbrace{e^{w \text{ after jump} - w \text{ before jump}}}_{\text{LRR}} - 1.$$

The price reaction comprises a dividend effect and a price-dividend ratio effect. If consumption changes by X upon the event, the dividend changes by $(1 + X)^\phi - 1$ due to leverage, reflecting the short-run risk. If $X < 0$, then $(1 + X)^\phi < 1$. Therefore, the exposure to normal jumps is negative and the risk premium for such jumps is positive. The price-dividend ratio effect, $e^{w \text{ after jump} - w \text{ before jump}}$, captures the long-run influence of such an event. It is smaller than 1 if and only if $e^{(\theta-1)(v \text{ after jump} - v \text{ before jump})}$ is greater than 1. Consequently, the resulting premium for long-run risk is always positive.

¹⁴A detailed derivation can be found in Appendix A.

Using the calibrations with a loss size of $L = -0.05$, Table 5 reports the decompositions of the risk premium for all model specifications. The peak-to-trough model solely allows for short-run risk leading to one jump risk premium only. Depending on the calibration, this premium is either 0.0098 or 0.0479. The significant difference can be attributed to the market price of risk. The amount of risk $\lambda[(1 + L)^\phi - 1]$ is equal to -0.0273 in the local calibration and -0.0105 in the disaster calibration. The market prices of risk are -0.3604 and -4.5739 , respectively. The disproportionately higher market price of risk in the disaster calibration overcompensates the smaller amount of risk, which leads to a larger equity risk premium. This finding mirrors the results of Barro (2006) and Rietz (1988) who propose disaster models to solve the equity premium puzzle. A risk-averse agent is more afraid of rare severe jumps as compared to frequent small jumps.

In the model with separated regime switches, there are two types of jump risk premia in every state: normal jumps and jumps triggering separated regime switches. Normal jumps denoted by the superscripts ‘ g, g ’ and ‘ b, b ’ lead to an immediate drop in consumption and dividends, but have no effect on the regime. Consequently, the market prices of risk $\eta^{g,g} = \eta^{b,b} = 1 - (1 + L)^{-\gamma} = -0.3604$ and the price exposures $\zeta^{g,g} = \zeta^{b,b} = (1 + L)^\phi - 1 = -0.0975$ only reflect short-run consumption risk. The risk premia $\lambda^{g,g}\eta^{g,g}\zeta^{g,g}$ and $\lambda^{b,b}\eta^{b,b}\zeta^{b,b}$ thus only differ because of different jump intensities. On the other hand, separated regime switches solely affect the economic regime, but have no immediate impact on consumption and dividends. The premium for regime switches from the good to the bad state is given by $\lambda^{g,b}\eta^{g,b}\zeta^{g,b}$. The market price of risk $\eta^{g,b} = 1 - e^{(\theta-1)(v^b-v^g)} = -1.4494$ reflects only long-run risk. Since the agent has a preference for early resolution of uncertainty, he demands a premium for the risk of adverse changes in the state variable. This implies that $\eta^{g,b}$ is negative. By a similar argument the market price $\eta^{b,g}$ for switches from the bad state to the good state is positive. The exposures are driven by the changes in the price-dividend ratios of the dividend claim. We have $\zeta^{g,b} = e^{w^b-w^g} - 1 = -0.2334$ and $\zeta^{b,g} = e^{w^g-w^b} - 1 = 0.3045$. Taken together, the agent demands a positive risk premium for both regime switches. It is equal to 0.0272 in the good state and 0.1304 in the bad state.

The model with jump-induced regime switches leads to the highest risk premium in the good state. The biggest share of this risk premium can be attributed to the premium for jump-induced regime switches. Since this regime switch happens simultaneously with a consumption drop, the market price of risk provides a combined compensation for short-

run and long-run risk:

$$\eta^{g,b} = 1 - \underbrace{(1+L)^{-\gamma}}_{1-\eta^{jump}} \underbrace{e^{(\theta-1)(v^b-v^g)}}_{1-\eta^{RS}} = \eta^{jump} + \eta^{RS} - \eta^{jump}\eta^{RS}. \quad (8)$$

The market price of risk is not just the sum of the market prices of risk for consumption jumps and regime switches, but involves an additional term $-\eta^{jump}\eta^{RS}$. In equilibrium, the representative agent demands this extra premium since the consumption drop sparks an economic crisis. Numerically, we get $\eta^{g,b} = -2.1449$, which can be decomposed into $\eta^{jump} = -0.3604$, $\eta^{RS} = -1.3118$, and $-\eta^{jump}\eta^{RS} = -0.4727$. The value of the interaction term is significant. In particular, its contribution to the market price of risk is larger than the contribution of consumption losses.¹⁵

To provide an intuition for the additional term in the market price of risk, we draw an analogy to the disaster model of Barro (2006) and Rietz (1988). In this model, the representative agent favors small jumps occurring with a large probability over one large jump happening with a small probability (given that the average loss is the same in both cases). The market price of risk and in turn also the equity premium in disaster models is thus nonlinear in the severity of the disasters. By the same line of argument, the agent dislikes joint downward jumps in consumption and in the wealth-consumption ratio much more than separate ones.

The exposure to jump-induced regime switches also combines a short-run and a long-run component:

$$\zeta^{g,b} = \underbrace{(1+L)^\phi}_{1+\zeta^{jump}} \underbrace{e^{w^b-w^g}}_{1+\zeta^{RS}} - 1 = \zeta^{jump} + \zeta^{RS} + \zeta^{jump}\zeta^{RS}. \quad (9)$$

The term ζ^{jump} is equal to -0.0975 and thus equal to $\zeta^{g,g}$. Moreover, we have $\zeta^{RS} = -0.2116$, and $\zeta^{jump}\zeta^{RS} = 0.0206$.¹⁶ Here the interaction term has a different sign than the other two terms. Having jump-induced regime switches instead of separated ones thus slightly reduces the overall exposure of the asset price to adverse jumps.

Finally, we turn to the risk premium for jump-induced regime switches which is the

¹⁵Notice that in the model with separated regime switches we have $\eta^{g,b} = -1.4494$. This is different from η^{RS} in the model with jump-induced regime switches because the wealth-consumption ratios differ in both economies.

¹⁶Note that, since the price-dividend ratios are in general different in the two models, ζ^{RS} is in general not equal to the exposure $\zeta^{g,b}$ in the model with separated regime switches.

combination of the market prices of risk and the exposure. It is given by

$$\begin{aligned} & \lambda^{g,b} \left[(1 + \zeta^{jump})(1 + \zeta^{RS}) - 1 \right] \left[1 - (1 - \eta^{jump})(1 - \eta^{RS}) \right] \\ & = \lambda^{g,b} \left(\zeta^{jump} + \zeta^{RS} + \zeta^{jump} \zeta^{RS} \right) \left(\eta^{jump} + \eta^{RS} - \eta^{jump} \eta^{RS} \right). \end{aligned}$$

Similar to the exposures and the market prices of risk, the risk premium is also superadditive, i.e. it is bigger than the sum of the single risk premia. Table 6 decomposes this risk premium into its components. The first part of the risk premium can be regarded as the compensation for (hypothetical) pure consumption jumps and (hypothetical) pure regime switches:

$$\lambda^{g,b} \zeta^{jump} \eta^{jump} + \lambda^{g,b} \zeta^{RS} \eta^{RS}.$$

This term is structurally equal to the risk premium in a model with separated regime switches.¹⁷ We get

$$\begin{aligned} \lambda^{g,b} \eta^{jump} \zeta^{jump} &= 0.0028 \\ \lambda^{g,b} \eta^{RS} \zeta^{RS} &= 0.0222. \end{aligned}$$

These two terms add up to around 0.025. In line with the long-run risk literature such as Bansal and Yaron (2004), the risk premium earned on long-run risk is much larger than the premium on short-run consumption risk.

The second part consists of cross risk premia. Since consumption jumps and regime switches happen simultaneously, consumption jumps also earn the market price of risk for regime switches, and vice versa:

$$\lambda^{g,b} \zeta^{jump} \eta^{RS} + \lambda^{g,b} \zeta^{RS} \eta^{jump}.$$

These two terms are unique to our model. Combining long-run and short-run risk thus increases the risk premium. This is so even without the super-additive terms in the exposures and market prices of risk. Numerically, the cross risk premia amount to

$$\begin{aligned} \lambda^{g,b} \eta^{RS} \zeta^{jump} &= 0.0102 \\ \lambda^{g,b} \eta^{jump} \zeta^{RS} &= 0.0061. \end{aligned}$$

They add up to around 0.016. Both are economically significant since both the market price of risk for long-run risk and the exposure to long-run risk are large.

¹⁷Numerically, it is slightly different since the changes in the price-dividend and wealth-consumption ratios upon a regime switch are different in the two models.

The third part of the risk premium are additional interaction terms resulting from super-additivity:

$$\lambda^{g,b} \zeta^{jump} \zeta^{RS} (\eta^{jump} + \eta^{RS}) - \lambda^{g,b} (\zeta^{jump} + \zeta^{RS}) \eta^{jump} \eta^{RS} - \lambda^{g,b} \zeta^{jump} \zeta^{RS} \eta^{jump} \eta^{RS}.$$

Numerically, they amount to

$$\begin{aligned} \lambda^{g,b} \zeta^{jump} \zeta^{RS} (\eta^{jump} + \eta^{RS}) &= -0.0028 \\ -\lambda^{g,b} (\zeta^{jump} + \zeta^{RS}) \eta^{jump} \eta^{RS} &= 0.0117 \\ -\lambda^{g,b} \zeta^{jump} \zeta^{RS} \eta^{jump} \eta^{RS} &= -0.0008 . \end{aligned}$$

They add up to another 0.008. Here, the most important contribution comes from $\eta^{jump} \eta^{RS}$, the super-additive component of the market prices of risk. Putting everything together, the overall premium for jump-induced regime switches is around 0.05. About half of this premium can be attributed to the joint occurrence of consumption jumps and regime switches.

To summarize, the additional risk premium in a model with jump-induced regime switches comes from two sources. First of all, consumption jumps inherit the market price of regime switches and regime switches inherit the market price of consumption jumps if these two events are tied together. Second, the superadditivity of exposures and market prices of risk adds further interaction terms to the risk premium.

The unconditional equity risk premium is also the largest in the economy with jump-induced regime switches. It exceeds the unconditional equity risk premium in a model with separated regime switches by around 1.5 percentage points although the local distribution of future consumption is the same in both economies. The link between long-run state variable risk and short-run consumption risk adds a significant component to the risk premium which is about the same order of magnitude as the premium for pure long-run risk or pure short-run risk.

4.5 Risk-free interest rate

Table 7 reports the risk-free rates in all economies. The unconditional risk-free rate in the model with jump-induced regime switches is 0.0155 and thus about 70 basis points lower than in the model with separated regime switches. In the peak-to-trough model, the numbers are 0.0284 (local calibration) and 0.0136 (disaster calibration). Similar to the analysis of the equity premium, the numbers show that our model with jump-induced

regime switches can also resolve the tradeoff between matching consumption data and generating a low risk-free rate.

The risk-free rate follows from the (negative) expected growth rate of the pricing kernel. It equals

$$r_f^j = \beta + \frac{1}{\psi} \frac{\mathbb{E}_t [dC_t]}{C_t dt} - \pi^{diff} - \sum_{k=g,b} \pi^{j,k},$$

in state $j \in \{g, b\}$. The risk-free interest rate comprises the subjective time preference rate β , the expected growth rate of consumption scaled by the inverse of the EIS, and several precautionary savings terms for the different risk factors in our model:

$$\begin{aligned} \pi^{diff} &= \frac{1}{2} \gamma \left(1 + \frac{1}{\psi}\right) \sigma^2 \\ \pi^{j,k} &= \lambda^{j,k} \left[-\eta^{j,k} + \frac{1}{\psi} L + \frac{1-\theta}{\theta} \left((1+L)^{1-\gamma} e^{\theta(v^k - v^j)} - 1 \right) \right]. \end{aligned}$$

for $(j, k) \in \{(g, g), (g, b), (b, b)\}$ and

$$\pi^{b,g} = \lambda^{b,g} \left[-\eta^{b,g} + \frac{1-\theta}{\theta} \left(e^{\theta(v^g - v^b)} - 1 \right) \right].$$

Table 7 also reports the decomposition of the risk-free rate into its components. The first three components (time preference rate, expected consumption growth, precautionary savings term for diffusive risk) are straightforward and identical across settings (except for the expected consumption growth in the disaster calibration).¹⁸

Since the good state prevails 90% of the time, differences in the unconditional risk-free rate are mainly driven by differences in the conditional risk-free rate in the good state. These differences arise from the precautionary savings terms for jump risk. In the economy with jump-induced regime switches, these terms reduce the risk-free rate by almost 1 percentage point as compared to the model with separated regime switches. The mechanism driving this finding is similar as for the equity premium. The precautionary savings term for jump-induced regime switches is $\pi^{g,b} = 0.0140$, compared to $\pi^{g,b} = 0.0053$ with separated regime switches.

¹⁸Notice that the expected consumption growth in the local calibration of the peak-to-trough model equals the unconditional expected consumption growth in the other two models. However, the expected consumption growth is slightly higher with the disaster calibration.

5 Robustness Checks

5.1 Elasticity of Intertemporal Substitution

The upper panel of Figure 2 depicts the local equity premium in the good state in the regime switching models and the unconditional local equity premium in the peak-to-trough specifications as functions of the elasticity of intertemporal substitution ψ . The other panels of Figure 2 depict the risk premium for regime switches from the good to the bad state, the market price of risk $\eta^{g,b}$ and the exposure $\zeta^{g,b}$ in the two regime switching models. The independent variable EIS ranges from 0.02 to 2 so that the figures capture our baseline case from above as well as the special case with CRRA preferences ($\psi = 1/6$) which is indicated by the left dotted line. Table 8 reports the same equity premium decomposition as Table 4, but for the power utility case ($\psi = 1/6$).

Comparing Table 8 to Table 4 confirms that only the risk premium for (jump-induced or separated) regime switches depends on the EIS. All other risk factors in our model (diffusion and normal jumps) are pure cash flow risk and do not affect the state variable. Therefore the exact degree of the agent's preference for early resolution of uncertainty (i.e. the tradeoff between γ and ψ) does not matter for these risk premia. The only part of the risk premium that changes upon a change in the EIS is the risk premium associated with regime switches. Consequently, the risk premium in the peak-to-trough model does not depend on ψ .

As Figure 2 shows, the total risk premia in the regime switching models are nonmonotonic functions of ψ with minima around the CRRA case ($\psi = 1/6$). This nonmonotonicity derives from the premium for (jump-induced or separated) regime switches $\lambda^{g,b}\eta^{g,b}\zeta^{g,b}$, which is depicted in the second panel. The lower graphs of Figure 2 depict $\eta^{g,b}$ and $\zeta^{g,b}$ as a function of the EIS. The nonmonotonic behavior of the risk premium comes from two ingredients: (i) the monotonicity of $\eta^{g,b}$ and $\zeta^{g,b}$ and (ii) their signs.

The market price $\eta^{g,b}$ is monotonically decreasing in ψ both for jump-induced and separated regime switches. Recall that we can write this market price as

$$\eta^{g,b} = 1 - \underbrace{(1+L)^{-\gamma}}_{1-\eta^{jump}} \underbrace{e^{(\theta-1)(v^b-v^g)}}_{1-\eta^{RS}}.$$

As the market price η^{jump} does not depend on ψ , the monotonic pattern comes from the market price for regime switches η^{RS} . For the log wealth-consumption ratios, it holds true that $v^b > v^g$ for $\psi < 1$ and $v^b < v^g$ for $\psi > 1$. In the latter case, the substitution effect

dominates the income effect and the wealth-consumption ratio is higher in the good state. The term $\theta - 1$ is negative for $\psi > 1$ and for $\psi < 1/6$ and positive in between. This results in a monotonically decreasing η^{RS} . For $\psi = 1/6$, η^{RS} is equal to zero, since the CRRA investor does not price state variable risk. For $\psi > 1/6$, adverse changes in the state command a negative market price of risk, which increases in absolute terms in the preference for early resolution of uncertainty (and thus in ψ).

In the model with separated regime switches, the market price $\eta^{g,b}$ is essentially equal to η^{RS} (except that the exact values for v^g and v^b are different). In the model with jump-induced regime switches, the market price of risk is amplified by the term involving η^{jump} . Therefore, $\eta^{g,b}$ is always smaller than the market price in the model with separated regime switches and has a null below $\psi = 1/\gamma$.

The exposures $\zeta^{g,b}$ are also decreasing in ψ . The more the investor cares about the timing of the resolution of uncertainty, the larger the price difference between the good and the bad state, and this price difference is reflected in the exposure $\zeta^{g,b}$. The exposure is always lower with jump-induced regime switches than with separated regime switches. This is because of the additional cash flow shock to which the regime switch is coupled. Similar to the market price of risk, the exposure switches sign and becomes positive if the EIS becomes sufficiently low for both model specifications. If the preference for early resolution of uncertainty is not very pronounced, a switch from the good to the bad regime increases all prices in economy. For CRRA utility ($\psi = 1/6$), we thus find the usual counterintuitive result that the price-dividend ratio is higher in bad states of the world than in good states. For recursive utility with $\psi > 1$, the direction of the relation reverses. Consequently, the exposure function has a null between $\psi = 1/6$ and $\psi = 1$.

The behavior of $\eta^{g,b}$ and $\zeta^{g,b}$ explains the nonmonotonicity of the jump risk premium depicted in the second graph. For high values of the EIS, both $\eta^{g,b}$ and $\zeta^{g,b}$ are negative and monotonically decreasing. Thus, their product is positive and monotonically increasing. For very low values of the EIS, both $\eta^{g,b}$ and $\zeta^{g,b}$ are positive and monotonically decreasing, which results in a positive and monotonically decreasing risk premium. In between, the risk premium has a minimum for $\psi \approx 0.2$. For values of ψ around $1/6$, the risk premium $\lambda^{g,b}\eta^{g,b}\zeta^{g,b}$ becomes even negative, so that the total risk premium in the first graph is close to zero.

Notice that the two curves in the first graph intersect around $\psi = 0.25$, i.e. for values below $\psi = 0.25$ the total risk premium is higher in a model with separated regime switches than with jump-induced regime switches. This is again in line with the well-known intuition

from the long-run risk literature that state variable risk and recursive preferences are necessary to generate sizeable risk premia. On the other hand, the difference in risk premia between both models is still sizeable for intermediate values of ψ between 0.5 and 1.5. Our findings are thus robust to changes of the EIS as long as the preference for early resolution of uncertainty is reasonably pronounced.

To summarize, the risk premium for regime switches reaches a minimum around $\psi = 1/\gamma$ which is the CRRA case. This minimum is even negative. For other values of ψ , the risk premium can be significantly larger. This pattern is particularly pronounced in an economy with jump-induced regime switches.

5.2 Jump Sizes and Intensities

The previous results have used the parametrization with a loss size of $L = -0.05$, i.e. Case III of Table 3. Figure 3 depicts the unconditional local equity premium and the risk-free rate for the other cases. The jump size L increases (in absolute terms) from Case I to Case V, while the jump intensity decreases. The figure reveals that the equity premium is high for all parametrizations of our model. It is the highest in Case IV. Overall, we conclude that there are several parameterizations of our model that can match the empirical consumption data and lead to a reasonable equity premium. In all cases, our additional channel of jump-induced regime switches increases the equity premium significantly. The additional premium is in all cases about the same order of magnitude as the premium for separated regime switches which is a pure long-run risk premium à la Bansal and Yaron (2004). The lower figure shows the risk-free rate for all cases. The effect of our new channel on the risk-free rate is essentially a mirror image of the findings for the equity premium.

5.3 Leverage

Figure 4 depicts the unconditional local equity premium as a function of the leverage parameter ϕ ranging between 1 and 3. So far, we have assumed $\phi = 2$. For large values of ϕ , the two models with regimes generate higher premia than both parametrizations of the peak-to-trough model. One reason for this finding is that the jump size in all our models is constant. In the peak-to-trough model, this implies that there is only one jump risk factor with a constant jump size, i.e. no jump *size* risk. In the regime switching models, conditional on the state of the economy, there are always two jump risk factors which have

different price impacts. Essentially, this is comparable to a situation with one jump risk factor where the jump size is drawn from a distribution with two possible outcomes. Stated differently, one can interpret the regime switching models, conditional on the economic state, as models with constant jump intensity, but jump *size* risk. The larger the leverage parameter, the larger is the average size and variance of price jumps. As compared to the peak-to-trough model, the risk premium in the regime switching models thus increases disproportionately as a function of the leverage parameter.

Finally, notice that the additional equity premium due to our jump-induced regime switch channel is robust to changes in the leverage parameter. This can be seen by comparing the red solid line and the blue dotted line. The difference shall be interpreted as this additional premium. Obviously, it is almost constant across leverage levels.

6 Conclusion

This paper entertains a tractable way to model the joint occurrence of negative consumption realizations and changes in the distribution of future consumption growth. This new channel combines short-run and long-run risk. We calibrate our model to consumption data covering 42 countries and show that it matches the data well. Our model generates disasters that are realistic both from their total size, but also from their duration. In particular, it avoids the so-called peak-to-trough calibration where the total size of a disaster is realized at a single point in time. By contrast, our model produces sufficiently many long-lasting periods in which moderate negative consumption realizations cluster. Nevertheless, this approach leads to a sizeable equity premium. Therefore, our results address the critique by Constantinides (2008). We also find that a stylized model that disentangles long-run and short-run risk cannot generate a realistic equity premium if the risk aversion of the representative agent is in a reasonable range.

Since our novel channel leads to a tractable class of models and allows for closed-form solutions of the equity premium and the risk-free rate, it could potentially be embedded in other asset pricing models. This might help future research to match the equity premium and the risk-free rate even if the focus is not particularly on this research question. In other words, our findings support the argument that by incorporating a combination of long-run and short-run risk one can achieve a sizeable equity premium without making unrealistic assumptions about the dynamics of crises.

A Solving for the Equilibrium

A.1 Wealth-Consumption Ratio

Let $Z_t \in \{g, b\}$ denote the state of the economy at time t . Then the representative investor has two value functions, one for each state:

$$J_t^{Z_t} = \mathbb{E}_t \left[\int_t^\infty f(C_s, J_s^{Z_s}) ds \right].$$

For the sake of readability, we will, however, suppress the dependence of the value function, the pricing kernel, the aggregate consumption and other variables on the state $Z_t \in \{good, bad\}$ in the following. As usual, the aggregator f is defined as

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

β denotes the subjective time discount rate, ψ the elasticity of intertemporal substitution, and γ the relative risk aversion. We also define $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$. A Feynman-Kac-like computation then gives

$$0 = f(C_t, J_t) + DJ_t \tag{10}$$

i.e. one Bellman equation for each state.

The dynamics of consumption in the good state are

$$\frac{dC_t}{C_{t-}} = \mu dt + \sigma dW_t + LdN_t^{g,g} + LdN_t^{g,b}, \tag{11}$$

its dynamics in the bad state are

$$\frac{dC_t}{C_{t-}} = \mu dt + \sigma dW_t + LdN_t^{b,b}. \tag{12}$$

We apply the following conjecture for the functional form of the value function J :

$$J = \frac{C^{1-\gamma}}{1-\gamma} \beta^\theta e^{\theta v^Z} \tag{13}$$

where v^Z can take two values, one in each state. Campbell, Chacko, Rodriguez, and Viceira (2004) and Benzoni, Collin-Dufresne, and Goldstein (2011) show that, with this conjecture, v^Z is the log wealth-consumption ratio. Plugging the guess (13) for J into the aggregator function results in

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

The infinitesimal generator DJ follows via Itô's Lemma:

$$\begin{aligned} DJ &= \left(1 - \frac{1}{\psi}\right) \theta J^g \mu - \frac{1}{2} \gamma \left(1 - \frac{1}{\psi}\right) \theta J^g \sigma^2 \\ &\quad + \lambda^{g,g} J^g \left[(1+L)^{1-\gamma} e^{\theta v^g - \theta v^g} - 1 \right] \\ &\quad + \lambda^{g,b} J^g \left[(1+L)^{1-\gamma} e^{\theta v^b - \theta v^g} - 1 \right] \end{aligned}$$

in the good state and

$$DJ = \left(1 - \frac{1}{\psi}\right) \theta J^b \mu - \frac{1}{2} \gamma \left(1 - \frac{1}{\psi}\right) \theta J^b \sigma^2 + \lambda^{b,b} J^b \left[(1+L)^{1-\gamma} e^{\theta v^b - \theta v^b} - 1\right] + \lambda^{b,g} J^b \left[e^{\theta v^g - \theta v^b} - 1\right]$$

in the bad state. Plugging these expressions into (10), dividing by θJ^g and θJ^b respectively, and rearranging some terms gives the following two algebraic equations for the two unknowns v^g and v^b :

$$\begin{aligned} 0 &= e^{-v^g} - \beta + \left(1 - \frac{1}{\psi}\right) \mu - \frac{1}{2} \gamma \left(1 - \frac{1}{\psi}\right) \sigma^2 \\ &\quad + \frac{1}{\theta} \lambda^{g,g} \left[(1+L)^{1-\gamma} - 1\right] + \frac{1}{\theta} \lambda^{g,b} \left[(1+L)^{1-\gamma} e^{\theta(v^b - v^g)} - 1\right] \\ 0 &= e^{-v^b} - \beta + \left(1 - \frac{1}{\psi}\right) \mu - \frac{1}{2} \gamma \left(1 - \frac{1}{\psi}\right) \sigma^2 \\ &\quad + \frac{1}{\theta} \lambda^{b,b} \left[(1+L)^{1-\gamma} - 1\right] + \frac{1}{\theta} \lambda^{b,g} \left[e^{\theta(v^g - v^b)} - 1\right]. \end{aligned}$$

A.2 Pricing Kernel

As Duffie and Epstein (1992a) and Benzoni, Collin-Dufresne, and Goldstein (2011) show, the pricing kernel is given by

$$\xi_t = \beta^\theta C_t^{-\gamma} e^{-\beta\theta t + (\theta-1) \left(\int_0^t e^{-v_u^Z} du + v_t^Z t \right)}. \quad (14)$$

The dynamics of the pricing kernel can be computed via Itô's Lemma. The partial derivatives of ξ with respect to C and v follow from (14). The dynamics of C are given in (11) and (12). The dynamics of the pricing kernel are

$$\begin{aligned} \frac{d\xi_t}{\xi_{t-}} &= \left[-\beta\theta + (\theta-1)e^{-v^g} \right] dt - \gamma\mu dt + \frac{1}{2}\gamma(1+\gamma)\sigma^2 dt \\ &\quad - \eta^{diff,g} dW_t - dN_t^{g,g} \eta^{g,g} - dN_t^{g,b} \eta^{g,b} \end{aligned}$$

in the good state and

$$\begin{aligned} \frac{d\xi_t}{\xi_{t-}} &= \left[-\beta\theta + (\theta-1)e^{-v^b} \right] dt - \gamma\mu dt + \frac{1}{2}\gamma(1+\gamma)\sigma^2 dt \\ &\quad - \eta^{diff,b} dW_t - dN_t^{b,b} \eta^{b,b} - dN_t^{b,g} \eta^{b,g} \end{aligned}$$

in the bad state. For later use, we abbreviate the drift of the pricing kernel by

$$\mu_\xi^Z = -\beta\theta + (\theta-1)e^{-v^Z} - \gamma\mu + \frac{1}{2}\gamma(1+\gamma)\sigma^2.$$

The market price of diffusion risk is $\eta^{diff,Z} = \gamma\sigma$. The market prices of jump risk are

$$\begin{aligned} \eta^{g,g} &= 1 - (1+L)^{-\gamma} \\ \eta^{g,b} &= 1 - (1+L)^{-\gamma} e^{(\theta-1)(v^b - v^g)} \\ \eta^{b,b} &= 1 - (1+L)^{-\gamma} \\ \eta^{b,g} &= 1 - e^{(\theta-1)(v^g - v^b)}. \end{aligned}$$

The risk-free rate is equal to the negative expected growth rate of the pricing kernel ξ_t :

$$\begin{aligned}
r_f^g &= \beta + \frac{1}{\psi} \left(\mu + L\lambda^{g,g} + L\lambda^{g,b} \right) - \frac{1}{2} \gamma \left(1 + \frac{1}{\psi} \right) \sigma^2 \\
&\quad - \lambda^{g,g} \left[-\eta^{g,g} + \frac{1}{\psi} L + \frac{1-\theta}{\theta} \left((1+L)^{1-\gamma} - 1 \right) \right] \\
&\quad - \lambda^{g,b} \left[-\eta^{g,b} + \frac{1}{\psi} L + \frac{1-\theta}{\theta} \left((1+L)^{1-\gamma} e^{\theta(v^b-v^g)} - 1 \right) \right] \\
r_f^b &= \beta + \frac{1}{\psi} \left(\mu + L\lambda^{b,b} \right) - \frac{1}{2} \gamma \left(1 + \frac{1}{\psi} \right) \sigma^2 \\
&\quad - \lambda^{b,b} \left[-\eta^{b,b} + \frac{1}{\psi} L + \frac{1-\theta}{\theta} \left((1+L)^{1-\gamma} - 1 \right) \right] \\
&\quad - \lambda^{b,g} \left[-\eta^{b,g} + \frac{1-\theta}{\theta} \left(e^{\theta(v^g-v^b)} - 1 \right) \right].
\end{aligned}$$

A.3 Pricing the Dividend Claim

For the price-dividend ratio of the claim to dividends, we apply the Feynman-Kac formula. Let w denote the log price-dividend ratio. Defining $g(\xi, D, w) = \xi D e^w$ results in

$$g(\xi_t, D_t, w_t) = \xi_t D_t e^{w_t} = \text{E}_t \left[\int_t^\infty \xi_\tau D_\tau d\tau \right] = \text{E}_t \left[\int_t^\infty \frac{g(\xi_\tau, D_\tau, w_\tau)}{e^{w_\tau}} d\tau \right].$$

The Feynman-Kac formula yields

$$\mathcal{D}g(\xi, D, w) + \frac{g(\xi, D, w)}{e^w} = 0 \quad \iff \quad \frac{\mathcal{D}g(\xi, D, w)}{g(\xi, D, w)} + e^{-w} = 0. \quad (15)$$

The dividend dynamics in our model are

$$\frac{dD_t}{D_{t-}} = \mu dt + \phi \sigma dW_t + [(1+L)^\phi - 1] dN_t^{g,g} + [(1+L)^\phi - 1] dN_t^{g,b}$$

in the good state and

$$\frac{dD_t}{D_{t-}} = \mu dt + \phi \sigma dW_t + [(1+L)^\phi - 1] dN_t^{b,b}$$

in the bad state. Itô's Lemma gives

$$\frac{\mathcal{D}g}{g} = \mu_\xi + \mu_D + \mu_w + \frac{1}{2} \frac{d\langle w^c \rangle}{dt} + \frac{d\langle \xi^c, D^c \rangle}{\xi D dt} + \frac{d\langle w^c, D^c \rangle}{D dt} + \frac{d\langle w^c, \xi^c \rangle}{\xi dt} + \text{Jump Terms.}$$

The log price-dividend ratio w can take the two values w^g and w^b only, i.e. μ_w is 0. From (15), we get the following two algebraic equations:

$$\begin{aligned}
0 &= e^{-w^g} + \mu_\xi^g + \mu - \eta^{diff,g} \phi \sigma \\
&\quad + \lambda^{g,g} [(1 - \eta^{g,g})(1 + L)^\phi - 1] \\
&\quad + \lambda^{g,b} [(1 - \eta^{g,b})(1 + L)^\phi e^{w^b - w^g} - 1] \\
0 &= e^{-w^b} + \mu_\xi^b + \mu - \eta^{diff,b} \phi \sigma \\
&\quad + \lambda^{b,b} [(1 - \eta^{b,b})(1 + L)^\phi - 1] \\
&\quad + \lambda^{b,g} [(1 - \eta^{b,g})e^{w^g - w^b} - 1].
\end{aligned}$$

A.4 Exposures and Equity Risk Premium

Conditional on the state, the dynamics of the asset price $P = e^w D$ follow via Itô's Lemma. In the good state, we have

$$\begin{aligned}
\frac{dP_t}{P_{t-}} &= \mu dt + \phi \sigma dW_t + [(1 + L)^\phi - 1] dN_t^{g,g} \\
&\quad + [(1 + L)^\phi e^{w^b - w^g} - 1] dN_t^{g,b}.
\end{aligned}$$

In the bad state, the dynamics are

$$\begin{aligned}
\frac{dP_t}{P_{t-}} &= \mu dt + \phi \sigma dW_t + [(1 + L)^\phi - 1] dN_t^{b,b} \\
&\quad + [e^{w^g - w^b} - 1] dN_t^{b,g}.
\end{aligned}$$

We abbreviate the sensitivities of the asset price to the different risk factors as

$$\begin{aligned}
\zeta^{diff} &= \phi \sigma \\
\zeta^{g,g} &= (1 + L)^\phi - 1 \\
\zeta^{g,b} &= (1 + L)^\phi e^{w^b - w^g} - 1 \\
\zeta^{b,b} &= (1 + L)^\phi - 1 \\
\zeta^{b,g} &= e^{w^g - w^b} - 1.
\end{aligned}$$

The expected excess return on the dividend claim, i.e. the equity risk premium, follows from these exposures and the respective market prices of risk. In the good state, it is equal to

$$\begin{aligned}
\gamma \phi \sigma^2 &+ \lambda^{g,g} [(1 + L)^\phi - 1] [1 - (1 + L)^{-\gamma}] \\
&+ \lambda^{g,b} [(1 + L)^\phi e^{w^b - w^g} - 1] [1 - (1 + L)^{-\gamma} e^{(\theta-1)(v^b - v^g)}].
\end{aligned}$$

The equity risk premium in the bad state is

$$\begin{aligned}
\gamma \phi \sigma^2 &+ \lambda^{b,b} [(1 + L)^\phi - 1] [1 - (1 + L)^{-\gamma}] \\
&+ \lambda^{b,g} [e^{w^g - w^b} - 1] [1 - e^{(\theta-1)(v^g - v^b)}].
\end{aligned}$$

References

- Backus, D., M. Chernov, and I. Martin, 2011, “Disasters Implied by Equity Index Options,” *Journal of Finance*, 66(6), 1969–2012.
- 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), 243–264.
- Barro, R., E. Nakamura, J. Steinsson, and J. Ursua, 2013, “Crises and Recoveries in an Empirical Model of Consumption Disasters,” *American Economic Journal: Macroeconomics*, 5(3), 35–74.
- Benzoni, L., P. Collin-Dufresne, and R. Goldstein, 2011, “Explaining Asset Pricing Puzzles Associated with the 1987 Market Crash,” *Journal of Financial Economics*, 101(3), 552–573.
- Campbell, J., G. Chacko, J. Rodriguez, and L. Viceira, 2004, “Strategic Asset Allocation in a Continuous-Time VAR Model,” *Journal of Economic Dynamics and Control*, 28(11), 2195–2214.
- Campbell, J., and J. Cochrane, 1999, “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107(2), 205–251.
- Chen, H., S. Joslin, and N.-K. Tran, 2012, “Rare Disasters and Risk Sharing,” *Review of Financial Studies*, 25, 2189–2224.
- Constantinides, G., 2008, “Comment on ‘Macroeconomic Crises since 1870’ by Barro and Ursua,” *Brookings Papers on Economic Activity*, pp. 341–350.
- Drechsler, I., and A. Yaron, 2011, “What’s Vol Got to Do with It?,” *Review of Financial Studies*, 24(1), 1–45.
- Duffie, D., and L. 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.
- Gabaix, X., 2012, “Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macrofinance,” *Quarterly Journal of Economics*, 127(2), 645–700.
- Goetzmann, W., and P. Jorion, 1999, “Global Stock Markets in the Twentieth Century,” *Journal of Finance*, 54(3), 953–980.
- Hung, M., 1994, “The Interaction between Nonexpected Utility and Asymmetric Market Fundamentals,” *Journal of Finance*, 49(1), 325–343.
- Julliard, C., and A. Ghosh, 2012, “Can Rare Events Explain the Equity Premium?,” *Review of Financial Studies*, 25(10), 3037–3076.
- Longstaff, F., and M. Piazzesi, 2004, “Corporate Earnings and the Equity Premium,” *Journal of Financial Economics*, 74(3), 401–421.
- Rietz, T., 1988, “The equity premium: A solution,” *Journal of Monetary Economics*, 22, 117–131.
- Tsai, J., and J. A. Wachter, 2014, “Rare booms and disasters in a multi-sector endowment economy,” Working Paper.
- Wachter, J., 2013, “Can time-varying risk of rare disasters explain aggregate stock market volatility?,” *Journal of Finance*, 68(3), 987–1035.

Panel A: Mean and standard deviation of consumption growth rates						
	annual mean	annual std	5-year mean	5-year std	10-year mean	10-year std
Data	0.020	0.065	0.103	0.182	0.221	0.280
Model: Case I	0.020	0.051	0.083	0.118	0.198	0.200
Model: Case II	0.020	0.054	0.082	0.128	0.195	0.218
Model: Case III	0.020	0.056	0.082	0.132	0.197	0.228
Model: Case IV	0.020	0.055	0.084	0.130	0.201	0.226
Model: Case V	0.020	0.052	0.086	0.122	0.205	0.214

Panel B: Tail probabilities								
	Percentage of country-year observations below ...							
	0	-0.02	-0.04	-0.05	-0.06	-0.08	-0.10	-0.15
Data	0.302	0.181	0.109	0.083	0.066	0.044	0.033	0.015
Model: Case I	0.312	0.189	0.111	0.084	0.064	0.038	0.023	0.005
Model: Case II	0.310	0.192	0.117	0.092	0.074	0.048	0.032	0.009
Model: Case III	0.306	0.190	0.117	0.093	0.075	0.051	0.035	0.013
Model: Case IV	0.302	0.185	0.112	0.088	0.071	0.047	0.032	0.012
Model: Case V	0.302	0.179	0.104	0.080	0.062	0.040	0.026	0.009

Table 1: Unconditional distribution of annual consumption growth rates

The table reports moments and tail probabilities of the unconditional distribution of consumption growth rates. The rows labeled ‘Data’ are taken from the Barro dataset and are discussed in Section 3.1. The model results have been obtained from Monte Carlo simulations over 500,000 years and are analyzed in Section 3.2. For brevity, the table reports results for the model with jump-induced regime switches only. The parameters for these cases are reported in Table 3.

Panel A: Threshold 0%

	1 year	2 years	3 years	4 years	5 years	6 years	7 years	8 years	9 years	10 years
Data	0.6160	0.2300	0.0920	0.0380	0.0140	0.0080	0	0.0020	0	0.0010
Model: Case I	0.6485	0.2186	0.0806	0.0316	0.0126	0.0049	0.0021	0.0007	0.0003	0.0001
Model: Case II	0.6484	0.2129	0.0817	0.0330	0.0136	0.0061	0.0023	0.0012	0.0004	0.0002
Model: Case III	0.6559	0.2074	0.0778	0.0325	0.0156	0.0060	0.0028	0.0014	0.0004	0.0002
Model: Case IV	0.6648	0.2050	0.0741	0.0312	0.0136	0.0061	0.0027	0.0015	0.0007	0.0002
Model: Case V	0.6782	0.2042	0.0702	0.0271	0.0116	0.0049	0.0022	0.0010	0.0006	0.0002

Panel B: Threshold -5%

	1 year	2 years	3 years	4 years	5 years	6 years	7 years	8 years	9 years	10 years
Data	0.7990	0.1520	0.0340	0.0120	0	0.0030	0	0	0	0
Model: Case I	0.8118	0.1488	0.0314	0.0064	0.0013	0.0001	0.0000	0.0000	0.0000	0.0000
Model: Case II	0.7702	0.1718	0.0421	0.0119	0.0030	0.0006	0.0001	0.0001	0.0000	0.0000
Model: Case III	0.7579	0.1726	0.0499	0.0139	0.0043	0.0011	0.0002	0.0001	0.0000	0.0000
Model: Case IV	0.7697	0.1632	0.0480	0.0132	0.0042	0.0012	0.0002	0.0001	0.0001	0.0000
Model: Case V	0.8097	0.1373	0.0388	0.0103	0.0029	0.0005	0.0003	0.0000	0.0001	0.0000

Panel C: Threshold -10%

	1 year	2 years	3 years	4 years	5 years	6 years	7 years	8 years	9 years	10 years
Data	0.8260	0.1290	0.0380	0.0080	0	0	0	0	0	0
Model: Case I	0.9099	0.0824	0.0071	0.0004	0.0002	0.0000	0.0000	0.0000	0.0000	0.0000
Model: Case II	0.8701	0.1137	0.0140	0.0019	0.0002	0.00001	0.0001	0.0000	0.0001	0.0001
Model: Case III	0.8482	0.1293	0.0188	0.0031	0.0005	0.0001	0.0000	0.0000	0.0000	0.0000
Model: Case IV	0.8507	0.1256	0.0203	0.0028	0.0006	0.0001	0.0000	0.0000	0.0000	0.0000
Model: Case V	0.8719	0.1111	0.0148	0.0019	0.0002	0.0000	0.0000	0.0000	0.0000	0.0000

Table 2: Distributions of the lengths of crisis periods

The table reports the unconditional distributions of the lengths of crises periods as discussed in Sections 3.1 and 3.2. The length of a crisis is defined as the number of consecutive years where the consumption growth rates are below a threshold (0% for Panel A, -5% for Panel B, -10% for Panel C). The table reports the unconditional probabilities that a crisis in the sample has a length of one up to ten years. The rows labeled 'Data' are based on the Barro dataset. The other rows are derived from Monte Carlo simulations over 500,000 years. For brevity, the table reports results for the model with jump-induced regime switches only. The parameters for these cases are reported in Table 3.

	Jump-induced regime switches					Separated regime switches				
	Case	Case	Case	Case	Case	Case	Case	Case	Case	Case
	I	II	III	IV	V	I	II	III	IV	V
μ	0.032	0.033	0.033	0.032	0.03	0.032	0.033	0.033	0.032	0.03
σ	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
L	-0.03	-0.04	-0.05	-0.06	-0.07	-0.03	-0.04	-0.05	-0.06	-0.07
$\lambda^{g,g}$	0.02	0.02	0.02	0.02	0.02	0.14	0.12	0.10	0.08	0.06
$\lambda^{g,b}$	0.12	0.1	0.08	0.06	0.04	0.12	0.1	0.08	0.06	0.04
$\lambda^{b,b}$	2.9	2.4	1.9	1.4	0.9	2.9	2.4	1.9	1.4	0.9
$\lambda^{b,g}$	1.08	0.9	0.72	0.54	0.36	1.08	0.9	0.72	0.54	0.36
$\lambda^{uncond.}$	0.416	0.348	0.28	0.212	0.144	0.416	0.348	0.28	0.212	0.144
p^{good}	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90
p^{bad}	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10
duration(good)	8.33	10	12.5	16.67	25	8.33	10	12.5	16.67	25
duration(bad)	0.93	1.11	1.39	1.85	2.78	0.93	1.11	1.39	1.85	2.78

	Peak-to-trough model									
	local calibration					disaster calibration				
	Case	Case	Case	Case	Case	Case	Case	Case	Case	Case
	I	II	III	IV	V	I	II	III	IV	V
μ	0.032	0.033	0.033	0.032	0.03	0.032	0.033	0.033	0.032	0.03
σ	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
L	-0.03	-0.04	-0.05	-0.06	-0.07	-0.220	-0.237	-0.249	-0.251	-0.245
λ	0.416	0.348	0.28	0.212	0.144	0.018	0.023	0.024	0.022	0.018

Table 3: Parameters of the consumption processes

The table reports the parameters of the consumption processes resulting from the calibration in Section 3. All parameters are annualized. The quantities below the horizontal line, $\lambda^{uncond.}$, ..., duration(bad), are derived from the parameters above the horizontal line.

	Diffusion risk	Pure cons. jumps	Jump-induced regime switches	Separated regime switches	Total premium
Model with jump-induced regime switches					
Good state	0.0192	0.0007	0.0495	—	0.0694
Bad state	0.0192	0.0668	—	0.1096	0.1956
Unconditional					0.0820
Model with separated regime switches					
Good state	0.0192	0.0035	—	0.0272	0.0499
Bad state	0.0192	0.0668	—	0.1304	0.2164
Unconditional					0.0665
Peak-to-trough model					
Local calibration	0.0192	0.0098	—	—	0.0290
Disaster calibration	0.0192	0.0479	—	—	0.0671

Table 4: Decomposition of the local equity risk premium

The table reports the local equity risk premium in the different economies as discussed in Section 4.4 as well as a decomposition into its various components. The jump risk premia in the second, third and fourth column are further decomposed in Table 5. The premium for jump-induced regime switches in the first row is also decomposed in Table 6. All results have been generated with the parameters from Case III reported in Table 3.

Model with jump-induced regime switches							
$\lambda^{g,g}$	0.02	$\lambda^{g,b}$	0.08	$\lambda^{b,b}$	1.9	$\lambda^{b,g}$	0.72
$\eta^{g,g}$	-0.3604	$\eta^{g,b}$	-2.1449	$\eta^{b,b}$	0.3604	$\eta^{b,g}$	0.5682
$\zeta^{g,g}$	-0.0975	$\zeta^{g,b}$	-0.288	$\zeta^{b,b}$	0.0975	$\zeta^{b,g}$	0.268
$RP^{g,g}$	0.0007	$RP^{g,b}$	0.0495	$RP^{b,b}$.0668	$RP^{b,g}$	0.1096
Model with separated regime switches							
$\lambda^{g,g}$	0.10	$\lambda^{g,b}$	0.08	$\lambda^{b,b}$	1.9	$\lambda^{b,g}$	0.72
$\eta^{g,g}$	-0.3604	$\eta^{g,b}$	-1.4494	$\eta^{b,b}$	-0.3604	$\eta^{b,g}$	0.5917
$\zeta^{g,g}$	-0.0975	$\zeta^{g,b}$	-0.2349	$\zeta^{b,b}$	-0.0975	$\zeta^{b,g}$	0.3070
$RP^{g,g}$	0.0035	$RP^{g,b}$	0.0272	$RP^{b,b}$	0.0668	$RP^{b,g}$	0.1304
Peak-to-trough model							
local	λ	0.28	disaster	λ		0.024	
calibration	η	-0.3604	calibration	η		-4.5739	
	ζ	-0.0975		ζ		-0.4360	
	RP	0.0098		RP		0.0479	

Table 5: Decomposition of the jump risk premia

The table decomposes all jump risk premia from Table 4 into jump intensities λ , market prices of risk η and price exposures ζ . All results have been generated with the parameters from Case III reported in Table 3.

Premium for pure consumption jumps	$\lambda^{g,b}\eta^{jump}\zeta^{jump}$	0.0028
Premium for pure regime switches	$\lambda^{g,b}\eta^{RS}\zeta^{RS}$	0.0222
Cross risk premia	$\lambda^{g,b}\eta^{RS}\zeta^{jump}$	0.0102
	$\lambda^{g,b}\eta^{jump}\zeta^{RS}$	0.0061
Super-additivity terms	$\lambda^{g,b}(\eta^{jump} + \eta^{RS})\zeta^{jump}\zeta^{RS}$	-0.0280
	$-\lambda^{g,b}\eta^{jump}\eta^{RS}(\zeta^{jump} + \zeta^{RS})$	0.0117
	$-\lambda^{g,b}\eta^{jump}\eta^{RS}\zeta^{jump}\zeta^{RS}$	-0.0008
Total	$\lambda^{g,b}\eta^{g,b}\zeta^{g,b}$	0.0495

Table 6: Decomposition of the risk premium for jump-induced regime switches

The table decomposes the risk premium for jump-induced regime switches given in the first row of Table 4 into its components as discussed in Section 4.4. All results have been generated with the parameters from Case III reported in Table 3.

	β	$\frac{1}{\psi} \frac{E_t[dC_t]}{C_t dt}$	Precautionary savings terms				Total risk-free rate
			Diffusion risk	Pure cons. jumps	Jump-ind. RS	Separated RS	
Model with jump-induced regime switches							
Good state	0.03	0.0140	-0.0072	-0.0003	-0.0140	—	0.0226
Bad state	0.03	-0.310	-0.0072	-0.0262	—	-0.0137	-0.0481
Unconditional							0.0155
Model with separated regime switches							
Good state	0.03	0.0140	-0.0072	-0.0014	—	-0.0053	0.0302
Bad state	0.03	-0.0310	-0.0072	-0.0262	—	-0.0152	-0.0496
Unconditional							0.0222
Peak-to-trough model							
Local calibration	0.03	0.0095	-0.0072	-0.0039	—	—	0.0284
Disaster calibration	0.03	0.0135	-0.0072	-0.0227	—	—	0.0136

Table 7: Decomposition of the local risk-free rate

The table reports the local risk-free rate in the different economies as discussed in Section 4.5 as well as a decomposition into its various components. All results have been generated with the parameters from Case III reported in Table 3.

	Diffusion risk	Pure cons. jumps	Jump-induced regime switches	Separated regime switches	Total premium
Model with jump-induced regime switches					
Good state	0.0192	0.0007	-0.0160	—	0.0039
Bad state	0.0192	0.0668	—	0	0.0860
Unconditional					0.0121
Model with separated regime switches					
Good state	0.0192	0.0035	—	0	0.0227
Bad state	0.0192	0.0668	—	0	0.0860
Unconditional					0.0290
Peak-to-trough model					
Local calibration	0.0192	0.0098	—	—	0.0290
Disaster calibration	0.0192	0.0479	—	—	0.0671

Table 8: Decomposition of the local equity premium (CRRA utility)

The table reports the local equity premium in the different economies as discussed in Section 5.1 as well as a decomposition into its various components. All results have been generated with the parameters from Case III reported in Table 3, but with an EIS of $\psi = 1/\gamma = 1/6$.




Figure 1: Histogram of peak-to-trough disaster sizes

The figure depicts the histogram of disaster sizes in our simulated data, measured by the peak-to-trough method. We simulate 500,000 years of consumption data using the model with jump-induced regime switches and the parameters from Case III. Next, we isolate all consumption-growth observations below 0%. Finally, if there are several consecutive observations below 0%, we compute the total drop in consumption over all these years and name this the peak-to-trough disaster size. The figures plot the histogram of these peak-to-trough disaster sizes.




Figure 2: Robustness check: elasticity of intertemporal substitution

The figure depicts the local equity premium in the good state (first picture), the risk premium for jump-induced and separated regime switches from the good to the bad state (second picture), the market price of risk $\eta^{g,b}$ (third picture) and the exposure $\zeta^{g,b}$ (fourth picture) as a function of ψ . The lines represent the model with jump-induced regime switches (red solid line), separated regime switches (blue dashed line) and the peak-to-trough model with local calibration (green dash-dotted line) and disaster calibration (black dash-dotted line). The vertical dotted lines mark $\psi = 1/\gamma = 1/6$ and $\psi = 1$. All results have been generated with the parameters from Case III reported in Table 3. All numbers are annualized.




Figure 3: Unconditional local equity premium and risk-free rate for all calibrations

The figure depicts the unconditional local equity premium (upper picture) and the unconditional local expected risk-free rate (lower picture) for all five cases reported in Table 3 (the numbers on the x-axis refer to the jump size). The lines represent the model with jump-induced regime switches (red solid line), separated regime switches (blue dashed line) and the peak-to-trough model with local calibration (green dash-dotted line) and disaster calibration (black dash-dotted line). All numbers are annualized.




Figure 4: Unconditional local equity premium as a function of the leverage parameter

The figure depicts the unconditional local equity premium as a function of ϕ . The lines represent the model with jump-induced regime switches (red solid line), separated regime switches (blue dashed line) and the peak-to-trough model with local calibration (green dash-dotted line) and disaster calibration (black dash-dotted line). All results have been generated with the parameters from Case III reported in Table 3. All numbers are annualized.

Recent Issues

No. 10	Tim Eisert, Christian Eufinger	Interbank network and bank bailouts: Insurance mechanism for non-insured creditors?
No. 9	Christian Eufinger, Andrej Gill	Basel III and CEO compensation in banks: Pay structures as a regulatory signal
No. 8	Ignazio Angeloni, Ester Faia, Marco Lo Duca	Monetary Policy and Risk Taking
No. 7	Matthieu Darraq Paries - Ester Faia - Diego Rodriguez Palenzuela	Bank and Sovereign Debt Risk Connection
No. 6	Holger Kraft - Eduardo Schwartz - Farina Weiss	Growth Options and Firm Valuation
No. 5	Grigory Vilkov - Yan Xiao	Option-Implied Information and Predictability of Extreme Returns
No. 4	Markku Kaustia - Elias Rantapuska	Does Mood Affect Trading Behavior?
No. 3	Markku Kaustia - Samuli Knüpfer - Sami Torstila	Stock Ownership and Political Behavior: Evidence from Demutualization
No. 2	Markku Kaustia - Elias Rantapuska	Does Mood Affect Trading Behavior?
No. 1	Dimitris Georgarakos - Michalis Haliassos - Giacomo Pasini	Household Debt and Social Interaction