Estimating the Dynamics of Consumption Growth

Gustavo Schwenkler*

June 19, 2018†

Abstract

We estimate models of consumption growth that allow for long-run risks and disasters using data for a series of countries over a time span of 200 years. Our estimates indicate that a model with small and frequent disasters that arrive at a mean-reverting rate best fits international consumption data. The implied posterior disaster intensity in such a model predicts equity returns without compromising the unpredictability of consumption growth. It also generates time-varying excess stock volatility, empirically validating key economic mechanisms often assumed in consumption-based asset pricing models.

*Schwenkler is at the Boston University Questrom School of Business, Department of Finance, 595 Commonwealth Avenue, Boston, MA 02215 (gas@bu.edu).

†We are grateful to Torben Andersen, Federico Bandi, Andrea Buffa, Christian Dorion, Philipp Illeditsch, Dino Palazzo, Eric Renault, Olivier Scaillet, Viktor Todorov, the seminar participants at Boston University, Brown University, Johns Hopkins University, Northwestern University, the European Central Bank, the University of Chicago, the 2017 Trends in Rio Conference, the 2017 INFORMS Annual Meeting, and the 2018 HEC-McGill Winter Finance Workshop for useful comments and suggestions.

1 Introduction

Some of the most successful asset pricing models explain asset pricing anomalies by postulating specific dynamics for aggregate consumption. Many times, the dynamics of consumption growth are not directly estimated from consumption data. Instead, they are calibrated to match asset pricing behavior. Such calibrations may be problematic if the target moments are not carefully chosen and if the model of consumption growth is misspecified. The formal estimation of consumption models with currently available methods is challenging due to the presence of latent factors and the low frequency nature of consumption data.[1]

In this paper, we estimate from consumption data for a series of countries models of aggregate consumption which nest long-run risk and disaster formulations that are popular in the literature. We accomplish this goal by exploiting the recently developed methodology of Guay and Schwenkler (2018) that enables maximum likelihood estimation, hypothesis testing, and comparative analysis of consumption models in the presence of latent factors and disasters. Our estimates deliver strong evidence in support of models that allow for time variation in the disaster rate as in the model of Wachter (2013). Such a specification is favored by the data over random walk formulations as in Campbell and Cochrane (1999), long-run risk formulations as in Bansal and Yaron (2004), or time-invariant disaster formulations as in Barro (2006). These results have important implications for the design of consumption-based asset pricing models.

We consider a continuous-time model in which the expected growth rate and the instantaneous volatility of consumption are driven by mean-reverting latent factors as in the long-run risk model of Bansal and Yaron (2004). At any point of time, consumption

[1] The consumption-based asset pricing literature recognizes these challenges. For example, Cochrane (2017) states: "There is some hope in formally testing models – do their moment conditions and cross-equation restrictions hold? – and in checking models' additional assumptions – do conditional moments vary as much and in the way that long-run risk or rare disaster models specify?" Similarly, Ludvigson (2013) claims: "Although an important first step, a complete assessment of leading consumption-based theories requires moving beyond calibration, to formal econometric estimation, hypothesis testing, and model comparison. Formal estimation, testing, and model comparison present some significant challenges, to which researchers have only recently turned."

can fall unexpectedly due to the arrival of a disaster as in the model of Barro (2006). The magnitude of a disaster is uncertain and the rate with with which disasters arrive varies over time in a mean-reverting way as in the model of Wachter (2013). We estimate our model using annual data on real per capita consumer expenditures for each of the countries of Australia, Germany, Japan, and the United States going back to 1834. We also estimate several nested submodels. The simplest model posits consumption as a random walk. It does not control for long-run risks or disasters. An intermediate model allows for persistent and transitory long-run risk components but not for disasters. Two additional models neglect long-run risks and instead control for disaster. The simpler disaster model only allows for a time-invariant disaster rate. The richer disaster model also allows for time variation in the disaster rate.

We pursue likelihood inference, which is feasible thanks to the novel methodology of Guay and Schwenkler (2018). With the methodology of Guay and Schwenkler (2018), we can compute parameter estimators that are consistent, asymptotically normal, and asymptotically efficient. This enables the use of standard t-tests to assess the significance of model parameters, as well as likelihood ratio tests for goodness-of-fit evaluations. The methodology of Guay and Schwenkler (2018) also delivers unbiased and computationally efficient estimators of posterior means and model-implied marginal distributions, which we use to assess the performance of our models. An alternative Bayesian approach is recently undertaken by Schorfheide et al. (2018) for the analysis of long-run risk models. We go beyond Schorfheide et al. (2018) by also controlling for disasters when evaluating long-run risk models. Furthermore, our likelihood-based approach makes it possible for us to perform a comparative analysis of the different nested models.

Our estimates indicate that time variation in the disaster rate is a predominant feature of international consumption data. Across countries, the average disaster rate, the disaster rate volatility, and the disaster rate persistence is estimated to be significantly large. Shocks to the disaster rate are estimated to be similarly persistent as in Wachter (2013). However, unlike Barro (2006), Barro and Jin (2011), Barro and Ursúa (2017), Wachter (2013) and others who find that disasters are extremely severe and extremely rare, we estimate that disasters are small and frequent: A disaster of a median size of

-0.75% arrives once every 10 months in the United States. We find similar estimates for Australia, Germany, and Japan. Our estimates of the frequency and magnitude of disasters are similar to those of Backus et al. (2011), who calibrate time-invariant disaster models to equity options data.

Our results show that the data favors a specification of consumption growth that allows for time-variation in the disaster rate. Likelihood ratio tests highlight that the introduction of time variation in the rate of disaster arrival always yields a significant improvement in the goodness-of-fit of a model. An analysis of the model-implied moments and marginal distributions of consumption growth across countries indicates that models that neglect the time variation of the disaster rate tend to overstate the center of the consumption growth distribution. Such overstatements can yield imprecise estimates of model parameters. For example, when we consider the nested model that only allows for persistent and transient long-run risk shocks, our parameter estimates are similar to those reported in the long-run risk literature by Bansal et al. (2012), Schorfheide et al. (2018), and others. However, when we consider the full model that also controls for time-varying disaster arrival, the estimates of the sensitivity of consumption growth to persistent and transitory long-run risk shocks are insignificant. These results suggest that the role of long-run risk shocks may be overstated in a model that does not control for time-varying disaster arrival. Similarly, we find that a model that only allows for time-invariant disaster arrival tends to understate the frequency of disasters. All in one, our results indicate that it is critical to control for time variation in the disaster rate to obtain accurate estimates of the parameters of consumption growth models.

We evaluate the implications of our results for macroeconomic and asset pricing modeling. The model-implied posterior sample paths of the expected consumption growth rate and the conditional consumption volatility are strongly time-varying and persistent due to the time-varying nature of the disaster rate. However, the persistent time variation of the expected consumption growth rate does not translate to consumption growth being predictable in the data because disasters are completely unpredictable events. These findings explain the seemingly contradictory results of Hansen et al. (2008), who uncovers evidence of persistent latent components in consumption growth by analyzing the long-run trade-off between risk and return, and those of Beeler and Campbell (2012), who cannot find any evidence of consumption growth predictability in U.S. data. The posterior disaster rate predicts the risk premium on the S&P 500 and is negatively associated with the 3-month Treasury Bill yield. Furthermore, the posterior disaster rate is positively related to the volatility of the S&P 500 but unrelated to the volatility of the 3-month Treasury Bill. These findings empirically validate key economic mechanisms in Wachter (2013) that explain important asset pricing phenomena.

The rest of this paper is organized as follows. Section 2 introduces our models. Section 3 summarizes our data and econometric approach. Section 4 presents our empirical results estimates. We analyze the implications of the time-varying disaster model in Section 5. Section 6 concludes. There are several technical appendices that summarize the methodology of Guay and Schwenkler (2018) and indicate how we apply it our setting.

2 A model of consumption growth

We posit a state-space model for consumption growth. In our model, the instantaneous growth rate of consumption fluctuates around its expected value due to the presence of several sources of risk. One natural source of risk is diffusive risk that would prevail if consumption growth behaved as a random walk. However, Schorfheide et al. (2018) recently reject the null hypothesis that consumption evolves as a random walk. We therefore allow for additional sources of risks that differentiate our model from the random walk model. Inspired by the long-run risk model of Bansal and Yaron (2004), we allow the expected consumption growth rate and the instantaneous consumption volatility to be dynamic and time-varying. Furthermore, we allow for the occurrence of unexpected drops in consumption as in the disaster literature (Barro (2006), Wachter (2013), and others). Overall, our model of consumption growth will feature four distinct sources of risk: diffusive risk, stochastic volatility, time-varying growth rate, and disastrous jumps.

We choose to work in continuous time when modeling consumption growth. We do so to exploit the recently developed methodology of Guay and Schwenkler (2018) for the exact and efficient estimation of multivariate continuous-time models with jumps via maximum likelihood. The availability of likelihood inference tools allows us to carry out a formal statistical analysis of the different sources of risk driving consumption growth in the data. If we were to model consumption growth in discrete time, the evaluation of the likelihood would be difficult. It would involve several approximations that may yield inefficient or biased parameter estimates.[2]

We lay out our model for consumption growth. Let a unit of time be one year and Δ be the frequency of the observation of consumption data (i.e., $\Delta = 1$ for annual data). We assume that $X_{1,t}$ is a stochastic process that measures log-consumption and evolves according to the stochastic differential equation

$$\mathrm{d}X_{1,t} = (\mu + g_t)\,\mathrm{d}t + \sigma v_t \mathrm{d}W_{1,t} + \mathrm{d}J_t, \quad X_{1,0} = 0 \tag{1}$$

for scalars $\mu \in \mathbb{R}$ and $\sigma > 0$. Here, W_1 is a standard Brownian motion, and

$$J_t = -\sum_{n\geq 1} \zeta E_n 1_{\{T_n \leq t\}}$$

is a pure-jump process with jump times $(T_n)_{n\geq 1}$ and stochastic intensity λ_t.

The factors g_t, v_t, and λ_t introduce time variation in the expected growth rate, the spot volatility, and the jump frequency of consumption. We specify the dynamics of these factors in the following sections.

2.1 Long-run risks

With v_t we model transitory long-run volatility shocks as in Bansal and Yaron (2004). The presence of v_t introduces stochastic volatility in consumption growth. We take $v_t = e^{vX_{3,t}}$ for a process X_3 that solves the stochastic differential equation

$$\mathrm{d}X_{3,t} = -\kappa_3 X_{3,t} + \mathrm{d}W_{3,t}. \tag{2}$$

with non-negative scalars v and κ_3. Here, W_3 is an independent standard Brownian motion, and $X_{3,0} = 0$ is fixed for simplicity.[3] The factor v_t scales the spot volatility of consumption growth up and down over time around the baseline value of σ.

[2] See Detemple et al. (2006), Giesecke and Schwenkler (2018), and Giesecke and Schwenkler (2017) for theoretical and numerical evidence on this issue.

[3] We have experiment with randomized initial value $X_{3,0}$ and found no major changes in the results.

The scalar κ_3 gives an inverse measure of the persistence of the stochastic volatility factor: We could rewrite (2) as an AR(1) model, in which case $e^{-\kappa_3 \Delta}$ would be the autoregressive coefficient. The scalar v, on the other hand, measures how sensitive consumption growth is to transitory long-run volatility shocks. When $v = 0$, the factor $v_t = 1$ almost surely for all $t \geq 0$ and consumption growth volatility is constant. In that case, there are no transitory long-run volatility shocks.

The factor g_t drives the expected growth rate of consumption. With this factor we model persistent long-run risk shocks as in Bansal and Yaron (2004). We assume that $g_t = \phi X_{2,t}$ for a stochastic process $X_{2,t}$ that solves the stochastic differential equation

$$\mathrm{d}X_{2,t} = -\kappa_2 X_{2,t} + v_t \mathrm{d}W_{2,t}, \tag{3}$$

where W_2 independent standard Brownian motion and $X_{2,0} = 0$.[4] The persistent long-run risk factor is centered around zero: $\mathbb{E}[g_t] = 0$ for all $t \geq 0$.[5] Because of this, the scalar μ in (1) measures the average consumption growth rate and ϕ measures the sensitivity of the expected consumption growth rate to long-run risk shocks. When $\phi = 0$, there are no persistent long-run risk shocks to the expected growth rate of consumption. When $\phi > 0$, the scalar κ_2 gives an inverse measure of the persistence of a long-run growth risk shock. Equation (3) can also be rewritten as an AR(1) model with $e^{-\kappa_2 \Delta}$ as the autoregressive coefficient.

2.2 Disasters

The jump process J in (1) introduces disasters as in the model of Barro (2006). Note that $J_t = J_s$ almost surely for any $0 \leq s < t < \infty$ unless a jump occurs at some point between times s and t. If a jump occurs at time T_n, then the process J jumps down, pulling log-consumption down with it: $\Delta X_{1,T_n} = \Delta J_{T_n} = -\zeta E_n$. Here, we assume that E_n is an independent sample of a standard exponentially distributed random variable and ζ is a non-negative scalar. The realization of $-\zeta E_n$ can be interpreted as a disaster in the sense of Barro (2006).

[4] Alternative fixed or random choices of $X_{2,0}$ have little impact on our results.

[5] However, the factor g does not have Gaussian conditional distributions due to the presence of the stochastic volatility factor v_t in Equation (3).

Disasters in our model arrive at the stochastic intensity λ_t that measures the conditional rate of disaster arrival. We assume that

$$\lambda_t = \ell_0 \exp\left(\ell_1 X_{4,t}\right)$$

for non-negative scalars ℓ_0 and ℓ_1. Here, X_4 is a stochastic process that satisfies

$$\mathrm{d}X_{4,t} = -\kappa_4 X_{4,t}\mathrm{d}t + \mathrm{d}W_{4,t}. \tag{4}$$

for an independent standard Brownian motion W_4, a non-negative scalar κ_4, and $X_{4,0} = 0$.[6] Because X_4 satisfies Ornstein-Uhlenbeck dynamics, its conditional law is Gaussian with finite variance. This implies that $\mathbb{E}[\int_0^\Delta \lambda_s \mathrm{d}s] < \infty$ so that the process J does not explode.

The process X_4 introduces time-variation in the arrival rate of disasters if $\ell_1 > 0$. In that case, the time-varying disaster intensity is mean-reverting, similarly as in the model of Wachter (2013). The arrival rate of disasters is time-invariant whenever $\ell_1 = 0$. The scalar κ_4 inversely measures the persistence of disaster intensity shocks: In an AR(1) formulation of (4), $e^{-\kappa_4\Delta}$ would be the autoregressive coefficient. The scalar ℓ_1 captures how strongly the disaster intensity varies over time around the baseline disaster intensity ℓ_0. In contrast, the scalar ζ measures the average magnitude of a disaster. There are no disasters whenever $\zeta = 0$ and $\ell_0 = 0$.

2.3 Nested models

In addition to the model specified in (1)–(4), we also analyze several models nested in our specification. We summarize the nested models in Table 1. A model that restricts $\phi = v = \zeta = 0$ has no long-run risk and no disasters. Such a model posits log-consumption as a random walk as in the model of Campbell and Cochrane (1999). We will estimate its free parameters μ and σ. An additional model we will analyze allows for persistent and transitory long-run risk shocks in the spirit of Bansal and Yaron (2004). Such a model restricts $\zeta = \ell_1 = 0$ and otherwise has 6 free parameters: μ, σ, ϕ, κ_2, v, and κ_3.[7] We also consider two types of disaster models. The simpler model assumes that the disaster

[6] Randomized choices of $X_{4,0}$ have no major influence on our results.

[7] An analogous model in discrete time is studied by Schorfheide et al. (2018).

intensity is constant as in Barro (2006). Such model restricts $\upsilon = \phi = \ell_1 = 0$ and has 4 free parameters: μ, σ, ζ, and ℓ_0. A more complex disaster model allows for a stochastic and mean-reverting disaster intensity as in Wachter (2013). This model has 6 free parameters $(\mu, \sigma, \zeta, \ell_0, \ell_1, \kappa_4)$ and restricts $\phi = \upsilon = 0$.

3 Data & estimation approach

Our goal is to estimate the vector $\theta = (\mu, \sigma, \phi, \kappa_2, \upsilon, \kappa_3, \zeta, \ell_0, \ell_1, \kappa_4)$ of parameters governing the dynamics of realized consumption growth, expected consumption growth, stochastic volatility, and disasters in (1)–(4). Each one of these parameters drives a different feature of the conditional distribution of consumption growth. We summarize the role played by each parameter in Table 2.

We use data on log-consumption to estimate the parameter vector θ via maximum likelihood. We construct time series of log-consumption from annual observations of real per capita consumer expenditure for a cross section of countries as recorded in the Barro-Ursúa data set.[8] We select the following countries for our analysis:

- Australia (available data: 1901-2009, 109 observations)
- Germany (available data: 1851-2009, 159 observations)
- Japan (available data: 1874-2009, 136 observations)
- United States (available data: 1834-2009, 176 observations)

For the U.S., we expand the Barro-Ursúa data by including data on real per capita consumer expenditure for the years 2010 through 2016, which we obtain from the Federal Reserve Bank of St. Louis FRED website. This gives us a total of 183 annual consumption observations for the United States. Table 3 provide summary statistics of our data.

Formally, we assume that the data is observed every Δ units of time and there are $m + 1$ data points available for inference. The total data available for inference is

$$\mathsf{D}_m = \{X_{1,0}, X_{1,\Delta}, \ldots, X_{1,m\Delta}\}.$$

[8] This data set can be downloaded from https://scholar.harvard.edu/barro/publications/barro-ursua-macroeconomic-data.

There is no data on the expected consumption growth factor g_t, the stochastic volatility factor v_t, and the disaster intensity factor λ_t. These processes are latent.

Because of the presence of latent factors, the likelihood is the projection of the complete-data likelihood ratio on the incomplete data space:[9]

$$\mathcal{L}_m(\theta) = \mathbb{E}_{\theta^*}\left[\prod_{i=1}^{m} \frac{p_\Delta\left(X_{(i-1)\Delta}, X_{i\Delta}; \theta\right)}{p_\Delta\left(X_{(i-1)\Delta}, X_{i\Delta}; \theta^*\right)} \,\middle|\, \mathsf{D}_m\right]. \tag{5}$$

Here, θ^* is the true data-generating parameter, $\mathbb{E}_\theta$ denotes the expectation operator under the measure $\mathbb{P}$ when the underlying parameter vector is θ, $X = (X_1, X_2, X_3, X_4)$ is the joint process defined in (1)–(4), and p_Δ is the transition density of the process X. A maximum likelihood estimator (MLE) $\hat{\theta}_m$ is an almost sure maximizer of the likelihood (5). We consider only MLE that satisfy the first order condition $\nabla\mathcal{L}_m(\hat{\theta}_m) = 0$.

The likelihood (5) poses a filter that is difficult to evaluate using standard filtering methods. We avoid these difficulties by using the recently developed simulation-based methodology of Guay and Schwenkler (2018). The methodology is computationally efficient and delivers accurate estimates of the likelihood $\mathcal{L}_m(\theta)$ and of MLE $\hat{\theta}_m$ with small computational costs. The resulting parameter estimators are consistent, asymptotically normal, and asymptotically efficient as the sample size m grows large. We provide details of the methodology and the asymptotic behavior of the estimators in Appendix A.

4 Model estimates

In this section, we present our model estimates and carry out a comparative analysis for the nested models. We show that a model that allows for time-varying disaster arrival is most supported by international consumption data.

4.1 Full model

Table 4 reports the parameter estimates of the full model for the different countries in our data sample. We see that the average growth rate and the average volatility parameters

[9] See Dembo and Zeitouni (1986) and Giesecke and Schwenkler (2018) for a formal derivation of the likelihood in incomplete data settings.

are significant across countries. We also see that the disaster intensity is estimated to be significantly positive, time-varying, and persistent in all countries. The disaster magnitude is significantly large in Australia and Germany. The parameters ϕ and v are not found to be significantly large in any country in our sample.

Across the board, we find that disasters are smaller and more frequent than calibrated in the disaster literature by Barro (2006), Barro and Jin (2011), and Barro and Ursúa (2017), among others. For example, Barro and Jin (2011) estimate that a disaster of median size of -19% arrives once every 26 years in the United States. In contrast, we estimate that in the U.S. a disaster of median size of -1.30% arrives once every 10 months.[10] Our estimates of the disaster magnitude and frequency for the U.S. are comparable to the estimates of Backus et al. (2011). Using options data, Backus et al. (2011) estimate that a disaster of median size -0.74% arrives once every 9 months.

Differences between our estimates and those of Barro (2006), Barro and Jin (2011), and Barro and Ursúa (2017) can be explained by differences in our econometric approaches. While we consider the whole distribution of consumption growth to estimate the disaster parameters in our likelihood-based approach, it is common in the disaster literature to calibrate these parameters to match the left-tail of realized consumption growth.[11] Such a constrained approach naturally leads to conservative disaster estimates. For example, if we were to calibrate the disaster parameters to match the left tail of U.S. consumption growth in our data set, we would find that a disaster of average size of -10% occurs once every 91 years. The left-tail-focused approach implies more severe disasters than those implied by the maximum likelihood estimates of Table 4.

We find that the disaster intensity is strongly time-varying and persistent in all countries in our data. These findings provide strong support for the time-varying disaster model of Wachter (2013). In the model of Wachter (2013), disasters arrive at an intensity that evolves as a Cox-Ingersoll-Ross process. In contrast, the disaster intensity in our model evolves as an exponential Ornstein-Uhlenbeck process. In spite of these differences,

[10]The long-run mean of the time-varying intensity in our model is $\ell_0 \exp(\frac{1}{2}\frac{\ell_1^2}{2\kappa_4})$ due to the Ornstein-Uhlenbeck dynamics of X_4 in (4).

[11]Typically, the disaster literature calibrates to match the distribution of realized consumption growth rates that exceed -10% in any given year.

our estimates of the persistence of disaster intensity shocks in the U.S. is similar to that estimated by Wachter (2013): Our estimate of κ_4 for the U.S. is 0.07 and the corresponding estimate of Wachter (2013) is 0.08. In contrast to Wachter (2013), however, our estimate of the baseline U.S. disaster intensity is much larger. Wachter (2013) estimates that a disaster occurs in the U.S. on average every 28 years while we find that a disaster occurs every 10 months. The differences between our disaster frequency estimate and that of Wachter (2013) can again be explained by differences in our econometric approaches. Wachter (2013) follows the standard approach in the disaster literature and calibrates the disaster frequency to match the left tail of consumption data. We instead follow a maximum likelihood approach that considers the whole distribution of consumption when estimating the disaster frequency.

Our estimates of the parameters ϕ and v are miniscule, which suggests that the latent factors g_t and v_t that model persistent and transitory long-run risk shocks in our model do not vary strongly over time. These findings lie in contrast to estimates by Bansal et al. (2012) and Schorfheide et al. (2018), who uncover significant latent process driving persistent and transitory long-run risk shocks using models that are nested in our specification. Unlike Bansal et al. (2012) and Schorfheide et al. (2018), we control for the occurrence of disasters when assessing the significance of the long-run risk parameters ϕ and v. As our analysis of the nested models in Section 4.2 shows, this difference between our approach and those of Bansal et al. (2012) and Schorfheide et al. (2018) explains the differences in our estimates.

4.2 Nested models

Tables 5 through 8 report the parameter estimates for the nested models fitted to Australian, German, Japanese, and U.S. consumption data. We see that the models that allow for time-invariant disasters (Model "DIS") cannot identify significant disasters in any country. Those models also understate the frequency of disaster arrival.

We identify significant disasters in Australia and Germany only when allowing for time-variation in the disaster intensity (Model "TVDIS"). For those countries, the parameter estimates for the average disaster magnitude ζ, the baseline disaster rate ℓ_0, the volatility ℓ_1 of the disaster intensity, and the speed of reversion κ_4 of the disaster intensity are significantly large and similar to the estimates reported for the full models in Table 4. We cannot find the average disaster magnitude ζ to be significantly large in Japan and the U.S., even though our estimates of the baseline disaster rate ℓ_0, the disaster intensity volatility ℓ_1, and the intensity speed of reversion κ_4 are found to be significantly large and similar to the the full model estimates of Table 4. For the U.S., our estimates for Model "TVDIS" suggest that a disaster of a median size of -0.75% arrives once every 10 months, almost precisely as often and severe as the disasters estimated by Backus et al. (2011) from options data. Overall, these results provide further support in favor of the time-varying disaster model of Wachter (2013), albeit with small and frequent disasters as in Backus et al. (2011).

The parameters ϕ and v of the model that allows for long-run risk shocks (Model "LRR") are estimated to be significantly large in Japan and the United States. In the U.S. in particular, our estimates are similar to the estimates of Bansal et al. (2012) and Schorfheide et al. (2018), who study discrete-time models analogous of Model "LRR". Bansal et al. (2012) calibrate their model to annual U.S. consumption and asset pricing data and estimate ϕ to be 0.0009, similar to our estimate of ϕ of 0.001. Schorfheide et al. (2018) estimate their model parameters via a Bayesian methodology using U.S. consumption data. Their findings imply the following 90% credible bands for our model parameters: $[0.001, 0.010]$ for ϕ, $[0.026, 0.055]$ for v, and $[0.012, 0.552]$ for κ_3. Our estimates of 0.001 for ϕ, 0.037 for v, and 0.017 for κ_3 lie inside of the credible bands implied by the results of Schorfheide et al. (2018).[12] All in one, when we consider a restricted model that allows for persistent and transitory long-run risk shocks but not for disasters, our estimates for the long-run risk parameters are comparable to those in the long-run risk literature. However, when we also control for the occurrence of disasters at a time-varying rate as in the full model, then the long-run risk parameters become small and insignificant (see Table 4). Our results suggest that models that do not control for time-varying disaster

[12] Our estimate of κ_2 for the U.S. is smaller that the estimates implied by the results of Bansal et al. (2012) and Schorfheide et al. (2018). We therefore estimate persistent long-run risk shocks in Model "LRR" to be more persistent than Bansal et al. (2012) and Schorfheide et al. (2018).

arrival may overstate the influence of long-run risk shocks. These findings explain why our estimates for the parameters ϕ and υ in the full model differ from the estimates Bansal et al. (2012) and Schorfheide et al. (2018) (see Section 4.1).

4.3 Likelihood ratio tests

Next, we run a comparative analysis of the nested models and the full model. We turn to likelihood ratio tests for this task.[13] Table 9 reports the likelihood ratio statistics for different model pairs. We see that the models that include long-run risk shocks ("LRR") or time-varying disasters ("TVDIS") are preferred over the simple random walk model ("RW") for all countries in our data. These results strongly reject the random walk hypothesis for consumption growth for a series of countries, complementing similar findings by Schorfheide et al. (2018) for the United States. The likelihood ratio tests do not provide support in favor of the time-invariant disaster model ("DIS"). Instead, they provide strong support in favor of the time-varying disaster model ("TVDIS"). We conclude that a formulation of disasters that arrive at a time-varying and mean-reverting rate as in Wachter (2013) is favored by the data over a time-invarying disaster formulation as in Barro (2006), Barro and Jin (2011), and Barro and Ursúa (2017).

The likelihood ratio tests of Table 9 also show that introducing time-varying disasters in a model that allows for long-run risk shocks (Models "FULL" vs "LRR") always leads to a significant increase in the goodness-of-fit of the model. In contrast, the inclusion of long-run risk shocks in a model with time-varying disasters (Models "FULL" vs. "TVDIS") does not lead to a significant increase in the fit to international consumption data. The estimates of the full and nested models reported in Tables 4–8 show that when we move from a restricted long-run risk model (Model "LRR") to the unrestricted full model, the estimates of the long-run risk parameters ϕ and υ shrink and become insignificant. These findings indicate that it is necessary to control for disasters that arrive at a time-varying rate when assessing the importance of long-run risk shocks. Ignoring the time-varying disaster channel may yield overstated estimates of long-run risk parameters.

[13] Likelihood ratio tests measure whether the better fit obtained by relaxing restricted parameters in a nested model is justified in light of the increased model complexity. We provide details in Appendix B.

All in one, the results of the likelihood ratio tests suggest that time-varying disasters are a prominent feature of consumption data in the countries we consider. Our results strongly favor models of consumption growth that allow for time-varying arrival of disasters as in Wachter (2013).

4.4 Implied moments and distributions

The model estimates of Tables 4–8 and the likelihood ratio tests of Table 9 indicate that models of consumption growth that do not control for time-varying disasters may overstate certain model components. To assess the impact of such overstatements, in Table 10 we report model-implied moments for the full and nested models for the countries in our data sample and compare to the realized moments in the data. We see that the models that control for long-run risks but not for disasters perform well at matching low-order and symmetric moments, such as the mean and the variance of consumption growth. In contrast, the models that control for disasters but not for long-run risks match the skewness and the left tail of the empirical consumption growth distribution. Consistent with the results of the likelihood ratio tests, we find that the model that controls for time-varying disasters (Model "TVDIS") performs best as it has the smallest moment-matching error is all countries. The full model performs worse than the model that ignores long-run risk shocks, suggesting that the full model is overparametrized.

We reach similar conclusions when looking at the model-implied marginal distributions of consumption growth displayed in Figures 1–4. Visually, the model that controls for time-varying disasters provides the best fit to the observed distribution of consumption growth in all countries. The models that neglect the time-varying disaster channel (Models "RW" and "LRR") tend to overstate the center of the empirical distribution. The model that controls both for long-run risks and time-varying disasters (Model "FULL") is overparametrized and fits the data less well than the model that only controls for time-varying disasters (Model "TVDIS").

Consistent with the results of the likelihood ratio test, our analysis of the model-implied moments and marginal distributions indicate that the time-varying disaster model best reflects the distribution of consumption growth in the data. A model that neglects the time-varying disaster channel will tend to overstate moments related to the center of the distribution. A model that incorporates sources of risks in excess of time-varying disasters will tend to be overparametrized, reducing its ability to fit the data.

5 Model implications

Our empirical results strongly support a model of consumption growth that allows for small disasters that arrive at a time-varying and mean-reverting rate. In this section, we study the macroeconomic and asset pricing implications of the time-varying disaster model specified in Table 1 (Model "TVDIS"). We focus on the United States as a sample. A key role will be played by the posterior sample path of the disaster intensity, $\mathbb{E}_t[\lambda_t]$, which is plotted in Figure 5. At any point of time, the posterior mean gives the best guess of the size of disaster intensity conditional on the available data at that time. We see in Figure 5 that the posterior mean of the disaster intensity fluctuates over time. There are periods of times in which the disaster intensity is close to zero and other times in which it reaches values of close to two disaster per year. The disaster intensity peaks especially high during the early 1900's, the First World War, and the Great Depression. Shocks to the disaster intensity are persistent, with a first-order autocorrelation coefficient of 0.57.

5.1 Conditional mean and variance of consumption growth

The time-varying disaster model posits that $g_t = v_t = 0$ almost surely for all $t \geq 0$ so that there are no transitory or persistent long-run risk shocks in the spirit of Bansal and Yaron (2004). Still, the fact that the disaster intensity fluctuates over time introduces time variation in the conditionally expected consumption growth rate and the conditional variance of consumption growth. Indeed, we have that the conditional expectation of the instantaneous growth rate of consumption is $\mu - \zeta \mathbb{E}_t[\lambda_t]$, where the last term accounts for the fact that the expected consumption growth rate will be low if many disasters are expected to arrive.[14] Similarly, the conditional variance of the instantaneous consumption growth rate is $\sigma^2 + \zeta^2 \mathbb{E}_t[\lambda_t]$, where the term $\zeta^2 \mathbb{E}_t[\lambda_t]$ accounts for the uncertainty in the

[14] More formally, the last term compensates for the jumps in (1).

magnitude and the frequency of disasters. Both the conditional expectation and the variance of the instantaneous consumption growth rate vary over time because the posterior mean of the disaster intensity ($\mathbb{E}_t[\lambda_t]$) varies over time (see Figure 5).

We plot the time series of the model-implied conditionally expected consumption growth rate and the conditional consumption growth volatility in Figure 6. We see that the model-implied conditional expectation and variance of consumption growth in the U.S. vary strongly over time, consistent with the empirical findings of Bansal et al. (2012). There were severely negative expected consumption growth rates during the early 1900's and the Great Depression. These periods were also associated with high consumption growth uncertainty. In Table 11 we regress the conditionally expected growth rate and the conditional volatility of consumption growth on several macroeconomic variables.[15] We see that periods of high expected consumption and low consumption volatility are associated with high inflation. We also see that the conditional sample paths of expected consumption growth and conditional consumption volatility are strongly autoregressive.

5.2 Long-run risks revisited

Figure 6 and the regression estimates of Table 11 highlight that the conditional expectation and volatility of consumption growth are persistent. In Table 13 we perform AR(1) regressions of expected consumption growth and conditional consumption volatility, and we estimate the autoregressive coefficients to be 0.570 and 0.587, respectively. The estimates of Table 13 imply monthly AR(1) coefficients of 0.954 for the expected consumption growth rate and 0.957 for the conditional consumption volatility. These monthly AR(1) coefficients are slightly smaller than those estimated by Schorfheide et al. (2018), which are 0.979 for expected consumption growth and 0.987 for the conditional volatility.

The results of Table 13 confirm that U.S. consumption data features persistent components in the expected consumption growth rate and the conditional consumption volatility, similarly as has been documented by Bansal and Yaron (2004), Bansal et al. (2010), Schorfheide et al. (2018), and others. Unlike the aforementioned papers, however, these dynamics arise in our model due to the persistent nature of disaster intensity shocks rather

[15] We provide summary statistics of the macroeconomic factors in Table 12.

than the presence of persistent latent factors in the mean and variance of consumption growth. These findings have important implications for the predictability of consumption. If there were persistent latent factors in the mean and the variance of consumption growth, then consumption growth itself would be predictable. However, consumption growth in our model is unpredictable because disasters are completely unanticipated events even when the disaster rate is predictable. In Table 14, we regress the realized consumption growth rate on lagged values of itself and of the expected consumption growth rate and the conditional consumption volatility. Consistent with the above discussion, we see that realized consumption growth rate is not predicted by the expected consumption growth rate or the conditional consumption volatility (or by itself).

The results of this section reconcile seemingly contradictory evidence by Hansen et al. (2008) and Beeler and Campbell (2012). By analyzing the long-run trade-off between risk and return of different asset classes, Hansen et al. (2008) uncovers evidence in support of the presence of persistent components in consumption growth. In contrast, Beeler and Campbell (2012) carry out a univariate statistical analysis of consumption growth in the U.S. and determine that there is little predictability in the data. These findinsg can coexist after recognizing that it is the time-varying nature of disaster arrival that introduce persistent components in the conditional mean and variance of consumption growth.

5.3 Asset pricing implications

The long-run risk model of Bansal and Yaron (2004) requires three components to match asset pricing data. First, there is a persistent latent component in the expected consumption growth rate. Second, the conditional consumption volatility exhibits persistent time-variation. Third, investors have recursive Epstein and Zin (1989) preferences. Several studies have highlighted the long-run risk model's ability to explain asset pricing phenomena; see Bansal et al. (2009), Bansal et al. (2010), and Hansen et al. (2008), among others. Given that our model features persistent expected growth rate and volatility components similar to those posited by the long-run risk literature, we conjecture that our time-varying disaster model should also be able to match asset pricing behavior as well as a long-run risk model.

To evaluate our conjecture, we check whether the expected consumption growth rate and the conditional consumption volatility of Figure 6 predict equity returns in the United States. For this, we regress the equity risk premium on the S&P 500 on lagged values of the expected consumption growth rate, the conditional consumption volatility, as well as several macroeconomic and financial factors that should drive risk premia. The results of these regressions can be found in Table 15. We see that a part of the equity risk premium on the S&P 500 is attributed to the the conditional volatility of consumption growth, even after controlling for the influence of volatility and macroeconomic factors. The regression results of Table 15 indicarte that shocks to the conditional consumption growth volatility carry a positive risk premium as in the model of Bansal and Yaron (2004).

Wachter (2013) integrates a time-varying disaster model similar to ours in an equilibrium model in which investors have recursive Epstein and Zin (1989) preferences. In the model of Wachter (2013), a key mechanism through which the time-varying disaster intensity has asset pricing implications is a precautionary savings effect: When the disaster intensity is high, the likelihood of large consumption drops is high so that investors have high desire to save. This implies that the risk-free rate should be low when the disaster intensity is high and vice versa. We test whether our time-varying disaster model implies similar dynamics for the risk-free rate. For this, in Table 16 we regress the yield of the 3-month Treasury Bill onto the posterior disaster intensity of Figure 5 and several macroeconomic and financial factors. We find that the coefficient on the posterior disaster intensity is significant and negative even when controlling for relevant explanatory factors. Our findings the precautionary savings mechanism of Wachter (2013) by showing that the 3-month Treasury Bill yield is indeed low when the posterior disaster intensity is large.

Wachter (2013) also posits that the volatility of stocks increases as the conditional volatility of consumption increases while the volatility of the risk-free rate remains unchanged. This mechanism of Wachter (2013) is important as it generates excess volatility that is reflected in equity prices but not in the risk-free rate, thus affecting risk premia. We verify that our model captures the positive association between stock and consumption volatility, and disassociation between risk-free rate and consumption volatility. In Table 17, we regress the annual realized volatility of the S&P 500 and the annual realized volatility of the 3-month Treasury bill yield on the posterior mean of the disaster intensity in Figure 5 and on several controls. We find that the volatility of the S&P 500 is high when the conditional volatility of consumption is high. We cannot find an association between the volatility of consumption growth and the volatility of the 3-month Treasury bill. These results empirically validate the mechanism posited by Wachter (2013).

All in one, our time-varying disaster model captures several mechanisms that are typically hardwired in general equilibrium consumption-based asset pricing models to match the data. Our model exhibits persistent fluctuations in the expected consumption growth rate and the conditional consumption volatility as in the model of Bansal and Yaron (2004). In the data, the aggregate volatility of stocks reacts positively to changes in our model-implied posterior mean of the disaster intensity, while the volatility of the risk-free rate is insensitive to changes in the posterior disaster intensity. Furthermore, the risk-free rate is negatively associated with the posterior mean of our disaster intensity. These results are consistent with mechanisms in Wachter (2013). Finally, the exposure to time-varying disaster risk is priced and carries a positive risk premium. In spite of our different formulation in which we assume no transitory and persistent long-run risk shocks and only allow for small disasters that arrive at a time-varying rate, the results of this section suggest that the model-implied posterior factors of our model match asset pricing behavior as well as the models of Bansal and Yaron (2004) and Wachter (2013).

6 Conclusion

In this paper, we estimate models of consumption growth that allow for time-varying arrival of disasters as well as persistent and transitory shocks in the the conditional mean and variance of consumption growth. Our model nests several formulations that are popular in the asset pricing literature, including the long-run risk model of Bansal and Yaron (2004), the disaster model of Barro (2006), and the time-varying disaster model of Wachter (2013). Our estimates strongly favor models of consumption that allow for small disasters that arrive at a time-varying and mean-reverting rate. The implied posterior means of the disaster intensity, the conditional expected growth rate, and the conditional consumption volatility in our model vary strongly over time in a persistent fashion, and explain asset pricing behavior. The results of this paper have important implications for the design of consumption-based asset pricing models.

A Likelihood estimation

We exploit the approach of Guay and Schwenkler (2018) for estimating the likelihood (5) and evaluating maximum likelihood estimators. There are two key challenges hindering the evaluation of the likelihood. First, the evaluation of the transition density p_Δ is challenging because the model specified in (1)-(4) is non-affine. Second, the likelihood (5) involves a conditional expectation that is evaluated with respect to the conditional distribution of the latent factors given the observed factor data. Such conditional expectations are difficult to evaluate for multivariate jump-diffusion model specified in (1)–(4). Guay and Schwenkler (2018) resolve the filtering problem by evaluating the likelihood under the null hypothesis that the true law generating the data is not $\mathbb{P}_{\theta^*}$ but rather an equivalent measure $\mathbb{P}^*$ under which the latent factors have simple dynamics and are independent of X_1. Indeed, Guay and Schwenkler (2018) show that

$$\mathcal{L}_m(\theta) \propto \mathbb{E}^* \left[\prod_{i=1}^m \frac{p_\Delta(X_{(i-1)\Delta}, X_{i\Delta}; \theta)}{\phi_l^*(X_{l,(i-1)\Delta}, X_{l,i\Delta})} \,\middle|\, \mathsf{D}_m \right], \tag{6}$$

where $\mathbb{E}^*$ denotes the expectation operator under the measure $\mathbb{P}^*$ and ϕ_l^* is the $\mathbb{P}^*$-transition density of the latent factors $X_l = (X_2, X_3, X_4)$.

Guay and Schwenkler (2018) also derive a simulation-based estimator of the transition density p_Δ. The estimator can be written as $\hat{p}_\Delta(v, w; \theta) = P(v, w; \theta, \mathsf{R})$ for an analytical function R that is known in algorithmic form and a vector R of $\mathcal{F}_0$-measurable random variables that do not depend upon the arguments $(v, w; \theta)$ of the density. The vector R contains simple random variables, such as standard uniforms, standard normals, and standard exponentials. Because of this, the density estimator has finite variance and can be computed in finite time. We can therefore replace the true density p_Δ with the simulated density $\hat{p}_\Delta$ without affecting the equality in (6).

All in one, Guay and Schwenkler (2018) evaluate an estimator of the likelihood (5) through the following steps:

(1) Generate samples $(\mathsf{R}^{(k)})_{k=1,\ldots,K}$ of the vector R of random variables,

(2) For $i = 0, \ldots, m$, generate samples $(\mathsf{X}_{l,i}^{(k)})_{k=1,\ldots,K}$ of the latent factor $X_{l,i\Delta}$ under the measure $\mathbb{P}^*$,

(3) Evaluate the Monte-Carlo estimator of (6) as

$$\hat{L}_m^K(\theta) = \frac{1}{K}\sum_{k=1}^{K}\prod_{i=1}^{m}\frac{P\left(\left(X_{1,(i-1)\Delta}, \mathsf{X}_{l,i-1}^{(k)}\right), \left(X_{1,i\Delta}, \mathsf{X}_{l,i}^{(k)}\right); \theta, \mathsf{R}^{(k)}\right)}{\phi_l^*\left(\mathsf{X}_{l,i-1}^{(k)}, \mathsf{X}_{l,i}^{(k)}\right)}. \tag{7}$$

Guay and Schwenkler (2018) show that the parameter $\hat{\theta}_m^K$ that maximizes the simulated likelihood (7) is consistent, asymptotically normal, and asymptotically efficient as $m \to \infty$ if $\frac{m}{K} \to 0$. A consistent estimator of the asymptotic variance-covariance matrix of $\hat{\theta}_m^K$ is

$$\left(\frac{1}{m}\sum_{i=0}^{m}\nabla \log\frac{\hat{L}_i^K(\hat{\theta}_m^K)}{\hat{L}_{i-1}^K(\hat{\theta}_m^K)}^\top \nabla \log\frac{\hat{L}_i^K(\hat{\theta}_m^K)}{\hat{L}_{i-1}^K(\hat{\theta}_m^K)}\right)^{-1}.$$

We specify the auxiliary measure $\mathbb{P}^*$ under which we evaluate the estimators of Guay and Schwenkler (2018). We assume that under $\mathbb{P}^*$:

- The sample paths $(X_{2,i\Delta})_{0\le i\le m}$, $(X_{3,i\Delta})_{0\le i\le m}$, and $(X_{4,i\Delta})_{0\le i\le m}$ of the latent factors are realizations of random walks that are independent of each other and of X_1, and
- The realizations $(X_{1,i\Delta} - X_{1,(i-1)\Delta})_{1\le i\le m}$ of consumption growth are i.i.d. normally distributed with mean μ^* and standard deviation σ^*.

These choices specify an equivalent measure given that $\mathbb{P}^*$ is nested in the family $\mathbb{P}_\theta$ by setting $\mu = \mu^*$, $\sigma = \sigma^*$, and $\phi = \upsilon = \zeta = \kappa_2 = \kappa_3 = \kappa_4 = 0$. For each country, we compute approximate likelihood estimators by maximizing the likelihood (7) with $K = 20,000$ and the measure $\mathbb{P}^*$ specified subject to the necessary parameter restrictions.

B Likelihood ratio tests

We use likelihood ratio tests to evaluate the incremental fit obtained by relaxing the restrictions of a nested model. Suppose that $\hat{\theta}_{m,R}^K$ and $\hat{\theta}_{m,U}^K$ are the parameter estimates computed with the methodology of Appendix A for a restricted model R and and unrestricted model U, respectively. The unrestricted model naturally achieves a better fit to the data than the restricted model because it has more parameters to match moments of the data. A fair evaluation of the incremental fit achieved by the more complex unrestricted model therefore needs to take into account the number of additional parameters used to achieve the better fit. An unrestricted model that achieves a significantly better fit with few additional parameters should be preferred over an unrestricted model that only achieves a marginally better fit but that has many additional parameters. Likelihood ratio tests provide a quantitative approach to evaluate the trade-off between better goodness-of-it and higher model complexity. Letting $\hat{L}_m^K$ denote the estimator (7) of the likelihood, standard econometric theory states that the likelihood ratio statistic implied by our methodology,

$$\mathrm{LR}(U,R) = 2\left(\log \hat{L}_m^K\left(\hat{\theta}_{m,U}^K\right) - \log \hat{L}_m^K\left(\hat{\theta}_{m,R}^K\right)\right), \tag{8}$$

has a chi-squared asymptotic distribution as $m, K \to \infty$ and $\frac{m}{K} \to 0$.[16] The degrees of freedom of the asymptotic chi-squared distribution is equal to the number of additional parameters in the unrestricted model. The availability of likelihood ratio test empowers us with tools to perform a formal comparative analysys of the full and nested models.

C Model-implied moments and distributions

Guay and Schwenkler (2018) derive estimators of model-implied moments and marginal distributions. Corollary 7.2. of Guay and Schwenkler (2018) shows that an unbiased estimator of the moment $\mathcal{E}_m(\theta; f, \theta') = \mathbb{E}_\theta\left[f(X_{t_1}, \ldots, X_{t_m}; \theta')\right]$ is given by

$$f\big((\mathsf{X}_{o,1}, \mathsf{X}_{l,1}), \ldots, (\mathsf{X}_{o,m}, \mathsf{X}_{l,m}); \theta'\big) \prod_{i=1}^m \frac{P_\Delta\big((\mathsf{X}_{o,i-1}, \mathsf{X}_{l,i-1}), (\mathsf{X}_{o,i}, \mathsf{X}_{l,i}); \theta \,\big|\, \mathsf{R}\big)}{\phi_o^*\left(\mathsf{X}_{o,i-1}, \mathsf{X}_{o,i}\right) \phi_l^*\left(\mathsf{X}_{l,i-1}, \mathsf{X}_{l,i}\right)},$$

where P_Δ is the density estimator of Guay and Schwenkler (2018) introduced in Appendix A, $(\mathsf{X}_{o,i})_{i=1,\ldots,m}$ is a $\mathbb{P}^*$-sample of the observable factor sample path $(X_{o,i\Delta})_{i=1,\ldots,m}$, and

[16] This result follows from the fact that the approximate parameter estimator $\hat{\theta}_m^K$ is consistent, asymptotically normal, and asymptotically efficient as $m \to \infty$ if $\frac{m}{K} \to 0$; see Guay and Schwenkler (2018).

$(\mathsf{X}_{l,i})_{i=1,\ldots,m}$ is a $\mathbb{P}^*$-sample of the latent factor sample path $(X_{l,i\Delta})_{i=1,\ldots,m}$. Here, $\mathbb{P}^*$ is an equivalent measure under which the latent and observable factors have simple dynamics and are independent of each other. A Monte Carlo estimator of the moment $\mathcal{E}_m(\theta; f, \theta')$ can therefore be constructed as follows:

(1) Generate samples $(\mathsf{R}^{(k)})_{k=1,\ldots,K}$ of the vector R of random variables,

(2) For $i = 0, \ldots, m$, generate samples $(\mathsf{X}_{l,i}^{(k)})_{k=1,\ldots,K}$ of the latent factor $X_{l,i\Delta}$ under the measure $\mathbb{P}^*$,

(3) For $i = 0, \ldots, m$, generate samples $(\mathsf{X}_{o,i}^{(k)})_{k=1,\ldots,K}$ of the observable factor $X_{o,i\Delta}$ under the measure $\mathbb{P}^*$,

(4) Compute the Monte Carlo estimator:

$$\frac{1}{K}\sum_{k=1}^{K} f\left(\left(\mathsf{X}_{o,1}^{(k)}, \mathsf{X}_{l,1}^{(k)}\right), \ldots, \left(\mathsf{X}_{o,m}^{(k)}, \mathsf{X}_{l,m}^{(k)}\right); \theta'\right) \prod_{i=1}^{m} \frac{P_\Delta\left(\left(\mathsf{X}_{o,i-1}^{(k)}, \mathsf{X}_{l,i-1}^{(k)}\right), \left(\mathsf{X}_{o,i}^{(k)}, \mathsf{X}_{l,i}^{(k)}\right); \theta \,\middle|\, \mathsf{R}\right)}{\phi_o^*\left(\mathsf{X}_{o,i-1}^{(k)}, \mathsf{X}_{o,i}^{(k)}\right) \phi_l^*\left(\mathsf{X}_{l,i-1}^{(k)}, \mathsf{X}_{l,i}^{(k)}\right)}. \tag{9}$$

The Monte Carlo estimator (9) is unbiased, computationally efficient, and converges to the true moment at square-root rate.

Similarly, Corollary 7.1 of Guay and Schwenkler (2018) shows that an unbiased estimator of the marginal density $p_{o,\Delta}(x, u; \theta)$ of the observable factor $X_{o,\Delta}$ is given by

$$\frac{P_\Delta\big((x, \mathsf{X}_{l,0}), (u, \mathsf{X}_{l,1}); \theta \,|\, \mathsf{R}\big)}{\phi_l^*\left(\mathsf{X}_{l,0}, \mathsf{X}_{l,1}\right)}.$$

We can therefore construct a Monte Carlo estimator of the marginal density $p_{o,\Delta}(x, u; \theta)$ as follows:

(1) Generate samples $(\mathsf{R}^{(k)})_{k=1,\ldots,K}$ of the vector R of random variables,

(2) For $i \in \{0, 1\}$, generate samples $(\mathsf{X}_{l,i}^{(k)})_{k=1,\ldots,K}$ of the latent factor $X_{l,i\Delta}$ under the measure $\mathbb{P}^*$,

(3) Compute the Monte Carlo estimator:

$$\frac{1}{K}\sum_{k=1}^{K} \frac{P_\Delta\big((x, \mathsf{X}_{l,0}^{(k)}), (u, \mathsf{X}_{l,1}^{(k)}); \theta \,|\, \mathsf{R}\big)}{\phi_l^*\left(\mathsf{X}_{l,0}^{(k)}, \mathsf{X}_{l,1}^{(k)}\right)}. \tag{10}$$

The Monte Carlo estimator (10) of the marginal density is also unbiased and computationally efficient. It converges to the true marginal density at square-root rate.

We specify our choice for the measure $\mathbb{P}^*$ under which we evaluate the moment and marginal density estimators of Guay and Schwenkler (2018). We assume that under $\mathbb{P}^*$:

- The sample path $(X_{j,i\Delta})_{0\leq i\leq m}$ of the latent factor X_j for $j \in \{2,3,4\}$ is a realization of an Ornstein-Uhlenbeck processes initialized at 0 with speed of reversion κ_j, long-term mean 0, and unit volatility,
- The latent factors X_2, X_3, and X_4 are independent of each other and of X_1, and
- The realizations $(X_{1,i\Delta} - X_{1,(i-1)\Delta})_{1\leq i\leq m}$ of consumption growth are i.i.d. normally distributed with mean μ and standard deviation σ.

These choices again specify an equivalent measure given that the measure $\mathbb{P}^*$ is nested in the family $\mathbb{P}_\theta$ by setting θ for which $\phi = v = \zeta = 0$. We evaluate the moment estimator (9) and the marginal density estimator (10) in Section 4.4 with $K = 10^6$ and the measure $\mathbb{P}^*$ defined above. We specify the function f necessary to evaluate the Monte Carlo estimator (9) for the different moments in Section 4.4 as follows:

- For the mean, we set $m = 1$ and $f((x_{o,1}, x_{l,1}); \theta') = x_{o,1} - X_{o,0}$.
- For the variance, we set $m = 1$ and $f((x_{o,1}, x_{l,1}); \theta') = (x_{o,1} - X_{o,0} - \hat{\mu})^2$, where $\hat{\mu}$ is the estimator of the mean.
- For the skewness, we set $m = 1$ and $f((x_{o,1}, x_{l,1}); \theta') = \left(\frac{x_{o,1}-X_{o,0}-\hat{\mu}}{\hat{\sigma}}\right)^3$, where $\hat{\mu}$ is the estimator of the mean and $\hat{\sigma}^2$ is the estimator of the variance.
- For the kurtosis, we set $m = 1$ and $f((x_{o,1}, x_{l,1}); \theta') = \left(\frac{x_{o,1}-X_{o,0}-\hat{\mu}}{\hat{\sigma}}\right)^4$, where $\hat{\mu}$ is the estimator of the mean and $\hat{\sigma}^2$ is the estimator of the variance.
- For the probability that the growth rate is less than x for $x \in \mathbb{R}$, we set $m = 1$ and $f((x_{o,1}, x_{l,1}); \theta') = 1_{\{x_{o,1}-X_{o,0}\leq x\}}$.
- For the first-order autocorrelation coefficient, we set $m = 2$ and $f((x_{o,1}, x_{l,1}), (x_{o,2}, x_{l,2}); \theta') = \frac{(x_{o,1}-X_{o,0}-\hat{\mu})(x_{o,2}-x_{o,1}-\hat{\mu})}{\hat{\sigma}^2}$, $\hat{\mu}$ is the estimator of the mean and $\hat{\sigma}^2$ is the estimator of the variance.

References

Backus, David, Mikhail Chernov and Ian Martin (2011), 'Disasters implied by equity index options', *The Journal of Finance* **66**(6), 1969–2012.

Bansal, Ravi and Amir Yaron (2004), 'Risks for the long run: A potential resolution of asset pricing puzzles', *The Journal of Finance* **59**(4), 1481–1509.

Bansal, Ravi, Dana Kiku and Amir Yaron (2010), 'Long run risks, the macroecon- omy, and asset prices', *Review of Financial Studies* pp. 542–546.

Bansal, Ravi, Dana Kiku and Amir Yaron (2012), 'An empirical evaluation of the long-run risks model for asset prices', *Critical Finance Review* **1**(1), 183–221.

Bansal, Ravi, Robert Dittmar and Dana Kiku (2009), 'Cointegration and consumption risks in asset returns', *Review of Financial Studies* **22**, 1343–1375.

Barro, Robert J. (2006), 'Rare disasters and asset markets in the twentieth century', *The Quarterly Journal of Economics* **121**(3), 823–866.

Barro, Robert J. and José F. Ursúa (2017), 'Stock-market crashes and depressions', *Research in Economics* **71**(3), 384 – 398.

Barro, Robert J. and Tao Jin (2011), 'On the size distribution of macroeconomic disasters', *Econometrica* **79**(5), 1567–1589.

Beeler, Jason and John Y. Campbell (2012), 'The long-run risks model and aggregate asset prices: An empirical assessment', *Critical Finance Review* **1**, 141–182.

Campbell, John Y. and John H. Cochrane (1999), 'By force of habit: A consumption?based explanation of aggregate stock market behavior', *Journal of Political Economy* **107**(2), 205–251.

Cochrane, John H. (2017), 'Macro-finance*', *Review of Finance* **21**(3), 945–985.

Dembo, A. and O. Zeitouni (1986), ‘Parameter estimation of partially observed continuous time stochastic processes via the EM Algorithm’, *Stochastic Processes and their Applications* **23**(1), 91 – 113.

Detemple, Jérôme, René Garcia and Marcel Rindisbacher (2006), ‘Asymptotic properties of monte carlo estimators of diffusion processes’, *Journal of Econometrics* **134**(1), 1–68.

Epstein, Lawrence and Stanley Zin (1989), ‘Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework’, *Econometrica* **57**, 937–969.

Giesecke, Kay and Gustavo Schwenkler (2017), Simulated Likelihood for Discretely Observed Jump-Diffusions. Working Paper.

Giesecke, Kay and Gustavo Schwenkler (2018), ‘Filtered likelihood for point processes’, *Journal of Econometrics* **204**, 33–53.

Guay, François and Gustavo Schwenkler (2018), Efficient inference and filtering for multivariate jump-diffusions. Working Paper.

Hansen, Lars Peter, John C. Heaton and Nan Li (2008), ‘Consumption strikes back? measuring long-run risk’, *Journal of Political Economy* **116**(2), 260–302.

Ludvigson, Sydney C. (2013), Advances in consumption-based asset pricing: Empirical tests, *in* ‘Handbook of the Economics of Finance’, Vol. 2, Part B, Elsevier, chapter 12, pp. 799 – 906.

Schorfheide, Frank, Dongho Song and Amir Yaron (2018), ‘Identifying long-run risks: A bayesian mixed-frequency approach’, *Econometrica* . Forthcoming.

Wachter, Jessica A. (2013), ‘Can time-varying risk of rare disasters explain aggregate stock market activity?’, *The Journal of Finance* **68**(3), 987–1035.

ModeLog-likelihood	Interpretation	Free parameters	Restricted parameters
RW	Random walk model	(μ, σ)	$\phi = \upsilon = \zeta = \ell_1 = 0$
LRR	Model that allows for transitory and persistent long-run risk components in the conditional growth rate and the conditional volatility of consumption growth. It is a continuous-time analogous of the model of Schorfheide et al. (2018).	$(\mu, \sigma, \phi, \kappa_2, \upsilon, \kappa_3)$	$\zeta = \ell_1 = 0$
DIS	Model that allows for disasters that arrive at a constant rate. This model is the continous-time analogous of the model of Barro (2006).	$(\mu, \sigma, \zeta, \ell_0)$	$\phi = \upsilon = \ell_1 = 0$
TVDIS	Model that allow for disasters that arrive at a time-varying and mean-reverting rate. This model is similar to the model of Wachter (2013), except that we model the disaster intensity as an exponential OU process.	$(\mu, \sigma, \zeta, \ell_0, \ell_1, \kappa_4)$	$\phi = \upsilon = 0$
FULL	Complete model that allows for persistent and transitory long-run risk components, and for the arrival of disasters at a time-varying and mean-reverting rate.	$(\mu, \sigma, \phi, \kappa_2, \upsilon, \kappa_3, \zeta, \ell_0, \ell_1, \kappa_4)$	

Table 1: This table summarizes the different models that we estimate in this paper.

Parameter	Interpretation	Analogy in the literature
μ	Annualized average consumption growth rate when no disaster occurs.	This parameter is the same as γ in Barro (2006), 12μ in Bansal and Yaron (2004), $12\mu_c$ in Schorfheide et al. (2018), and μ in Wachter (2013).
σ	Annualized average consumption growth volatility in the absence of stochastic volatility and disasters.	This parameter is the same as σ in Barro (2006) and Wachter (2013), and analogous to $\sqrt{12}\sigma$ in Bansal and Yaron (2004) and Schorfheide et al. (2018).
ϕ	Sensitivity of the expected consumption growth rate to persistent long-run risk shocks.	This parameter is equivalent to the parameter $12\psi_e$ in the long-run risk model of Bansal and Yaron (2004), or $12\psi_x\sigma$ in Schorfheide et al. (2018).
κ_2	Speed of reversion of the persistent latent factor driving the expected consumption growth rate.	This parameter is inversely related the persistence of long-run growth rate shocks. It can be translated to a monthly AR(1) coefficient similar to the parameter ρ in the long-run risk model of Bansal and Yaron (2004) and Schorfheide et al. (2018) by computing $e^{-\kappa_2/12}$.
v	Degree of time-variation of the conditional consumption volatility.	This parameter is the same as $\sqrt{12}\sigma_{h_c}$ in Schorfheide et al. (2018).
κ_3	Speed of reversion of stochastic consumption volatility.	This parameter can be translated to a monthly AR(1) coefficient $e^{-\kappa_3/12}$ that is analogous to the parameter ν_1 in Bansal and Yaron (2004) or ρ_{h_c} in Schorfheide et al. (2018).
ζ	Average disaster magnitude.	This parameter is equivalent to the contraction size b in Barro (2006).
ℓ_0	Baseline disaster frequency.	This parameter is equivalent to the baseline disaster frequency p in Barro (2006) and $\bar{\lambda}$ in Wachter (2013).
ℓ_1	Degree of time-variation of the disaster intensity.	This parameter is equivalent to the parameter σ_λ of Wachter (2013).
κ_4	Speed of reversion of the time-varying disaster intensity.	This parameter is analogous to the parameter κ in Wachter (2013).

Table 2: *Parameters*. This table summarizes the roles played by each of the parameters in Model (1)-(4).

	AU	DE	JP	US
Mean	0.014	0.015	0.022	0.015
Standard deviation	0.050	0.038	0.069	0.038
Skewness	−0.982	−0.071	−1.509	−0.071
Kurtosis	7.910	3.528	19.887	3.528
Median	0.017	0.018	0.021	0.016
Minimum	−0.221	−0.205	−0.433	−0.100
Maximum	0.170	0.202	0.357	0.110
Probability of a decline of more than 5%	0.074	0.038	0.059	0.038
Probability of a decline of more than 10%	0.028	0.005	0.022	0.005
First-order autocovariance	0.000	0.000	0.001	0.000
Number of observations	109	159	136	176
Time span	1901 – 2009	1851 – 2009	1874 – 2009	1834 – 2016

Table 3: Summary statistics of our data. We construct time series of log-consumption from annual observations of real per capita consumer expenditure recorded in the Barro-Ursúa data set, which can be downloaded from `https://scholar.harvard.edu/barro/publications/barro-ursua-macroeconomic-data`. We extend the time series for the U.S. with data from the Federal Reserve Bank of St. Louis FRED website

	AU	DE	JP	US
μ	** 0.0194	*** 0.0245	*** 0.0307	*** 0.0129
	(0.0066)	(0.0057)	(0.0035)	(0.0013)
σ	*** 0.0525	*** 0.0572	*** 0.0755	*** 0.0356
	(0.0048)	(0.0036)	(0.0038)	(0.0013)
ϕ	3×10^{-10}	1×10^{-9}	9×10^{-6}	8×10^{-10}
	(0.0022)	(0.0015)	(0.0012)	(0.0005)
κ_2	0.0477	* 0.0470	** 0.0622	* 0.0097
	(0.0272)	(0.0250)	(0.0221)	(0.0048)
v	0.0008	0.0036	3×10^{-8}	6×10^{-10}
	(0.0088)	(0.0057)	(0.0047)	(0.0049)
κ_3	0.0123	0.0082	0.0036	0.0099
	(0.0107)	(0.0080)	(0.0084)	(0.0062)
ζ	* 0.0466	*** 0.0157	0.0393	0.0433
	(0.0209)	(0.0036)	(0.0484)	(0.0407)
ℓ_0	*** 1.7759	*** 0.7593	*** 0.8865	*** 0.3546
	(0.2687)	(0.1134)	(0.1658)	(0.0508)
ℓ_1	*** 0.5637	*** 0.4487	*** 0.4217	*** 0.4932
	(0.1028)	(0.0456)	(0.0561)	(0.0313)
κ_4	* 0.0867	** 0.0918	** 0.0844	** 0.0703
	(0.0395)	(0.0342)	(0.0286)	(0.0249)
Log-likelihood	352.72	487.00	372.89	627.40

Table 4: Parameter estimates of the full model for Australia ("AU"), Germany ("DE"), Japan ("JP"), and the United States ("US"). We compute these estimates by maximizing the simulated likelihood of Appendix A. We use the Nelder-Mead algorithm in R to compute the maximizer of the simulated likelihood. Given in parenthesis are asymptotic standard errors. All estimates are measured in annual terms. *** indicates significance at the 99.9% confidence level, ** indicates significance at the 99% confidence level, and * indicates significance at the 95% confidence level.

	AU			
	RW	LRR	DIS	TVDIS
μ	*** 0.013	*** 0.016	*** 0.013	*** 0.020
	(0.002)	(0.004)	(0.002)	(0.005)
σ	*** 0.050	*** 0.041	*** 0.050	*** 0.052
	(0.002)	(0.003)	(0.002)	(0.002)
ϕ		4×10^{-10}		
		(0.001)		
κ_2		0.004		
		(0.007)		
υ		*** 0.051		
		(0.006)		
κ_3		$\cdot$ 0.017		
		(0.010)		
ζ			0.035	* 0.012
			(2435921)	(0.006)
ℓ_0			0.000	*** 1.757
			(0.096)	(0.265)
ℓ_1				*** 0.564
				(0.090)
κ_4				* 0.084
				(0.039)
Log-likelihood	275.19	285.67	275.19	349.94

Table 5: Parameter estimates of the nested models for Australia ("AU"). We compute these estimates by maximizing the simulated likelihood of Appendix A subject to the restrictions in Table 1. We use the Nelder-Mead algorithm in R to compute the maximizer of the simulated likelihood. Given in parenthesis are asymptotic standard errors. All estimates are measured in annual terms. *** indicates significance at the 99.9% confidence level, ** indicates significance at the 99% confidence level, * indicates significance at the 95% confidence level, and $\cdot$ indicates significance at the 90% confidence level.

	DE			
	RW	LRR	DIS	TVDIS
μ	*** 0.021	*** 0.017	*** 0.021	*** 0.024
	(0.003)	(0.002)	(0.003)	(0.004)
σ	*** 0.052	*** 0.037	*** 0.052	*** 0.055
	(0.002)	(0.001)	(0.002)	(0.002)
ϕ		3×10^{-4}		
		(3×10^{-4})		
κ_2		*** 0.023		
		(0.006)		
v		*** 0.126		
		(0.009)		
κ_3		** 0.054		
		(0.016)		
ζ			0.010	** 0.038
			(554993)	(0.011)
ℓ_0			0.000	*** 0.758
			(0.080)	(0.096)
ℓ_1				*** 0.448
				(0.038)
κ_4				** 0.092
				(0.035)
Log-likelihood	395.99	422.98	395.99	484.35

Table 6: Parameter estimates of the nested models for Germany ("DE"). We compute these estimates by maximizing the simulated likelihood of Appendix A subject to the restrictions in Table 1. We use the Nelder-Mead algorithm in R to compute the maximizer of the simulated likelihood. Given in parenthesis are asymptotic standard errors. All estimates are measured in annual terms. *** indicates significance at the 99.9% confidence level, ** indicates significance at the 99% confidence level, * indicates significance at the 95% confidence level, and · indicates significance at the 90% confidence level.

	JP			
	RW	LRR	DIS	TVDIS
μ	*** 0.023	*** 0.019	*** 0.025	*** 0.032
	(0.006)	(0.003)	(0.006)	(0.003)
σ	*** 0.075	*** 0.039	*** 0.075	*** 0.076
	(0.002)	(0.001)	(0.002)	(0.001)
ϕ		* 0.001		
		(0.000)		
κ_2		* 0.015		
		(0.011)		
υ		*** 0.138		
		(0.006)		
κ_3		** 0.058		
		(0.024)		
ζ			0.000	0.001
			(19278322)	(0.021)
ℓ_0			0.000	*** 0.891
			(0.086)	(0.179)
ℓ_1				*** 0.421
				(0.061)
κ_4				** 0.084
				(0.029)
Log-likelihood	299.53	336.44	299.49	369.76

Table 7: Parameter estimates of the nested models for Japan ("JP"). We compute these estimates by maximizing the simulated likelihood of Appendix A subject to the restrictions in Table 1. We use the Nelder-Mead algorithm in R to compute the maximizer of the simulated likelihood. Given in parenthesis are asymptotic standard errors. All estimates are measured in annual terms. *** indicates significance at the 99.9% confidence level, ** indicates significance at the 99% confidence level, * indicates significance at the 95% confidence level, and · indicates significance at the 90% confidence level.

	US			
	RW	LRR	DIS	TVDIS
μ	*** 0.015	*** 0.012	*** 0.015	*** 0.013
	(0.001)	(0.001)	(0.001)	(0.001)
σ	*** 0.038	*** 0.034	*** 0.038	*** 0.036
	(0.002)	(0.002)	(0.002)	(0.001)
ϕ		** 0.001		
		(0.000)		
κ_2		0.005		
		(0.009)		
v		*** 0.037		
		(0.005)		
κ_3		*** 0.017		
		(0.004)		
ζ			0.294	0.025
			(14131627)	(0.058)
ℓ_0			3×10^{-8}	*** 0.356
			(0.074)	(0.049)
ℓ_1				*** 0.491
				(0.029)
κ_4				** 0.071
				(0.025)
Log-likelihood	554.02	562.18	554.02	626.60

Table 8: Parameter estimates of the nested models for the United States ("US"). We compute these estimates by maximizing the simulated likelihood of Appendix A subject to the restrictions in Table 1. We use the Nelder-Mead algorithm in R to compute the maximizer of the simulated likelihood. Given in parenthesis are asymptotic standard errors. All estimates are measured in annual terms. *** indicates significance at the 99.9% confidence level, ** indicates significance at the 99% confidence level, * indicates significance at the 95% confidence level, and $\cdot$ indicates significance at the 90% confidence level.

Models		AU	DE	JP	US
LRR vs. RW	Statistic	*** 20.96	*** 54.60	*** 73.94	*** 16.68
	Degrees of freedom	4	4	4	4
	p-value	0.000	0.000	0.000	0.076
TVDIS vs. RW	Statistic	*** 149.50	*** 176.74	*** 139.64	*** 145.16
	Degrees of freedom	4	4	4	4
	p-value	0.000	0.000	0.000	0.000
FULL vs. TVDIS	Statistic	5.56	5.28	6.34	1.60
	Degrees of freedom	4	4	4	4
	p-value	0.235	0.260	0.175	0.809
FULL vs. LRR	Statistic	*** 134.10	*** 127.42	*** 72.04	*** 130.08
	Degrees of freedom	4	4	4	4
	p-value	00.00	0.000	0.000	0.000

Table 9: Likelihood ratio tests of the null hypothesis that a constrained model fits the data as well as an unconstrained model. The test statistic is given by two-times the logarithm of the ratio between the maximum likelihood of the unconstrained and the constrained models. The asymptotic distribution of the test statistic is chi-square distribution with degrees of freedom equal to the number of parameters that the unconstrained model has in excess of the constrained model. We provide details in Appendix B. *** indicates significance at the 99.9% confidence level.

AU					
	Data	RW	LRR	TVDIS	FULL
Mean	0.014	0.013	0.016	−0.005	−0.028
Standard deviation	0.050	0.055	0.045	0.050	0.071
Skewness	−0.982	0.022	0.022	0.435	−0.131
Kurtosis	7.910	3.000	2.914	3.535	2.937
Probability of a decline of more than 5%	0.074	0.129	0.070	0.164	0.377
Probability of a decline of more than 10%	0.028	0.020	0.005	0.018	0.144
Relative RMSE		0.639	0.640	0.577	0.634

DE					
	Data	RW	LRR	TVDIS	FULL
Mean	0.017	0.020	0.017	0.004	0.013
Standard deviation	0.053	0.058	0.041	0.070	0.064
Skewness	−0.550	0.022	−0.012	−0.225	0.084
Kurtosis	7.264	3.000	2.931	3.006	2.832
Probability of a decline of more than 5%	0.057	0.115	0.055	0.211	0.167
Probability of a decline of more than 10%	0.038	0.021	0.001	0.075	0.030
Relative RMSE		0.600	0.597	0.587	0.615

JP					
	Data	RW	LRR	TVDIS	FULL
Mean	0.022	0.023	0.019	0.029	0.009
Standard deviation	0.069	0.084	0.044	0.078	0.087
Skewness	−1.509	0.006	0.018	0.124	0.038
Kurtosis	19.887	2.960	2.961	3.214	2.939
Probability of a decline of more than 5%	0.059	0.190	0.060	0.142	0.244
Probability of a decline of more than 10%	0.022	0.073	0.002	0.048	0.103
Relative RMSE		0.852	0.852	0.840	0.853

US					
	Data	RW	LRR	TVDIS	FULL
Mean	0.015	0.015	0.011	0.007	0.003
Standard deviation	0.038	0.038	0.037	0.042	0.048
Skewness	−0.071	0.000	0.003	−0.211	−0.695
Kurtosis	3.528	3.000	2.913	3.358	4.056
Probability of a decline of more than 5%	0.038	0.044	0.051	0.084	0.115
Probability of a decline of more than 10%	0.005	0.001	0.001	0.011	0.038
Relative RMSE		0.151	0.175	0.064	0.233

Table 10: Model-implied moments for the full and nested models. We compute model-implied moments using the methodology of Guay and Schwenkler (2018); details are given in Appendix C. We compare the model-implied moments to the corresponding moments in the data. All moments are measured in annual terms. The relative RMSE is the ratio of the root mean squared error across all moments for a given country over the Euclidean norm of the moments in the data for that country.

	Expected growth rate	Conditional standard deviation
Constant	−0.001	*** 0.019
	(0.003)	(0.004)
Real GDP growth	−0.007	0.002
	(0.019)	(0.006)
Inflation rate	* 0.053	* −0.015
	(0.022)	(0.006)
S&P 500 return	0.028	−0.007
	(0.055)	(0.016)
S&P 500 realized volatility	−0.004	0.001
	(0.033)	(0.010)
Recession indicator	0.000	0.000
	(0.003)	(0.001)
Geopolitical risk	0.000	0.000
	(0.000)	(0.000)
Lagged value	*** 0.508	*** 0.527
	(0.095)	(0.094)
Adjusted R^2	0.35	0.37
Number of observations	86	86
Time span	1930–2016	1930–2016

Table 11: Ordinary least-squares regressions of the conditionally expected consumption growth rate ($\mu - \zeta\mathbb{E}_t[\lambda_t]$, Column "Expected growth rate") and the conditional consumption standard deviation ($\sqrt{\sigma^2 + \zeta^2\mathbb{E}_t[\lambda_t]}$, Column "Conditional standard deviation") on several macroeconomic factors and their lagged values. Here, we use the sample paths in Figure 6 for the expected growth rate and the conditional standard deviation of consumption growth. We provide summary statistics of the regressors in Table 12. The sampling interval is annual. For factors that are sampled at monthly frequencies, we aggregate to annual frequencies by taking annual averages. *** indicates significance at the 99.9% confidence level and * indicates significance at the 95% confidence level.

	Recession indicator	Real GDP growth	Inflation
Mean	0.296	0.033	0.033
Standard deviation	0.456	0.049	0.050
Skewness	0.896	0.011	0.551
Kurtosis	−1.197	3.155	2.944
Median	0	0.033	0.027
Minimum	0	−0.129	−0.158
Maximum	1	0.189	0.237
Number of observations	1957	87	1236
Time span	Dec 1854–Dec 2016	1930–2016	Jan 1914–Dec 2016
Sampling frequency	Monthly	Annual	Monthly

	S&P 500 return	S&P 500 volatility	3M T-Bill Yield
Mean	0.006	0.046	0.035
Standard deviation	0.054	0.028	0.032
Skewness	0.290	2.971	1.031
Kurtosis	9.564	3.155	1.072
Median	0.009	0.040	0.030
Minimum	−0.299	0.011	0.0001
Maximum	0.422	0.223	0.163
Number of observations	1091	1069	996
Time span	Feb 1926–Dec 2016	Dec 1926–Dec 2016	Jan 1934–Dec 2016
Sampling frequency	Monthly	Monthly	Monthly

	3M T-Bill volatility	Geopolitical risk	Equity risk premium
Mean	0.002	78.061	0.049
Standard deviation	0.003	68.824	0.130
Skewness	3.619	2.513	−0.610
Kurtosis	15.714	8.957	0.676
Median	0.001	55.415	0.063
Minimum	0	4.950	−0.338
Maximum	0.020	622.330	0.373
Number of observations	984	1416	90
Time span	Jan 1935–Dec 2016	Jan 1899–Dec 2016	1927–2016
Sampling frequency	Monthly	Monthly	Annual

Table 12: Summary statistics of our control variables. We compute a recession indicator using the NBER recession dates. Data on real GDP growth, inflation, the S&P 500 index, and the 3-month Treasury bill are obtained from the Federal Reserve Bank of St. Louis FRED website. Here, inflation is measured as the percentage change from a year ago in the non seasonally adjusted CPI for all urban consumers. We compute the realized volatility of the S&P 500 as the trailing volatility over the past 12 monthly returns. The volatility of the 3-month Treasury bill is computed as the trailing volatility of the past 12 increments of the yield. S&P returns and volatility, as well as the T-Bill volatility are measured on a monthly scale. We obtain a geopolitical risk indicator from Dario Caldara and Matteo Iacoviello website (https://www2.bc.edu/matteo-iacoviello/gpr.htm#overview). This is a sentiment-based index the measures how often words related to geopolitical threats are mentioned in the New York Times, Chicago Tribune, and the Washington Post. We measure the equity risk premium as the unlevered aggregate equity risk premium on the S&P 500 by taking the difference between the real return on the S&P 500 index and the real return on the 3-month T-Bill, divided by 1.5 to account for leverage.

	Expected growth rate	Conditional standard deviation
Constant	* -0.001	*** 0.017
	(0.001)	(0.002)
AR(1)	*** 0.570	*** 0.587
	(0.061)	(0.060)
Adjusted R^2	0.32	0.34
Number of observations	182	182
Time span	$1835 - 2016$	$1835 - 2016$

Table 13: AR(1) regressions of the conditionally expected consumption growth rate ($\mu - \zeta\mathbb{E}_t[\lambda_t]$, Column "Expected growth rate") and the conditional consumption standard deviation ($\sqrt{\sigma^2 + \zeta^2\mathbb{E}_t[\lambda_t]}$, Column "Conditional standard deviation"). We use the sample paths in Figure 6 for the expected growth rate and the conditional volatility of consumption growth. The sampling frequency of the data is annual. *** indicates significance at the 99.9% confidence level, and ** indicates significance at the 99% confidence level.

	Realized consumption growth	Realized consumption growth
Constant	*** 0.015	−1.016
	(0.003)	(1.166)
Lagged realized consumption growth	0.041	0.043
	(0.074)	(0.074)
Lagged expected consumption growth		8.033
		(8.804)
Lagged conditional consumption variance		25.971
		(29.379)
Adjusted R^2	−0.004	−0.007
Number of observations	181	181
Time span	1837 – 2016	1837 – 2016

Table 14: Regressions of the realized consumption growth rate on lagged values of itself and of the expected consumption growth rate and the conditional consumption volatility. We use the sample paths in Figure 6 for the expected growth rate and the conditional volatility of consumption growth. The sampling frequency of the data is annual. *** indicates significance at the 99.9% confidence level.

Lagged values of:	ERP	ERP	ERP
S&P volatility	* 0.438	** −1.339	*** −0.950
	(0.077)	(0.138)	(0.141)
Expected consumption growth		−0.402	−0.355
		(0.470)	(0.447)
Conditional consumption volatility		*** 2.720	*** 2.302
		(0.183)	(0.258)
Inflation			** −0.259
			(0.096)
Real GDP growth			*** 1.082
			(0.079)
Geopolitical risk			*** -4×10^{-4}
			(7×10^{-5})
Adjusted R^2	0.03	0.20	0.32
Number of observations	1069	1069	1032
Time span	Jan 1927–Jan 2016	Jan 1927–Jan 2016	Feb 1930–Jan 2016

Table 15: Regressions of the equity risk premium on the posterior means of the expected consumption growth and the conditional consumption volatility, as well as control factors. We construct a time series of the unlevered aggregate equity risk premium in the U.S. by taking the difference between the real return on the S&P 500 index and the real return on the 3-month T-Bill, divided by 1.5 to account for leverage; see Table 12 for summary statistics. All data is sampled at a monthly frequency and the regressors are lagged by one month. For controls that are observed at an annual frequency, we construct monthly observations by interpolating with constants. We do not include an intercept in the regression because the equity risk premium should only respond to active sources of risk. *** indicates significance at the 99.9% confidence level, ** indicates significance at the 99% confidence level, and * indicates significance at the 95% confidence level.

	T-Bill	T-Bill
Constant	*** 0.050	*** 0.036
	(0.002)	(0.002)
Posterior disaster intensity	*** −0.024	*** −0.017
	(0.003)	(0.003)
Inflation		*** 0.396
		(0.026)
Real GDP growth		*** −0.077
		(0.021)
S&P 500 return		0.004
		(0.019)
Geopolitical risk		0.000
		(0.000)
Adjusted R^2	0.06	0.25
Number of observations	997	997
Time span	Jan 1934–Dec 2016	Jan 1934–Dec 2016

Table 16: Regressions of the 3-Month Treasury Bill yield on the posterior mean of the disaster intensity and several control factors. All data is sampled at a monthly frequency, and the regressors are contemporaneous. The posterior disaster intensity is the sample path of $\mathbb{E}_t[\lambda_t]$ displayed in Figure 5. *** indicates significance at the 99.9% confidence level.

	S&P 500 volatility	S&P 500 volatility	S&P 500 volatility	T-Bill volatility	T-Bill volatility
Constant	*** 0.051	*** 0.042	*** 0.049	0.000	0.000
	(0.003)	(0.006)	(0.006)	(0.000)	(0.000)
Lagged posterior disaster intensity		· 0.013	* 0.018		−0.001
		(0.008)	(0.008)		(0.007)
Lagged S&P 500 return	*** −0.715	*** −0.693	*** −0.613		
	(0.172)	(0.171)	(0.162)		
Lagged T-Bill volatility				*** 0.512	*** 0.505
				(0.092)	(0.092)
Lagged T-Bill yield				*** 0.035	*** 0.034
				(0.008)	(0.008)
Lagged real GDP growth			−0.188		
			(0.052)		
Lagged geopolitical risk			0.000		
			(0.000)		
Adjusted R^2	0.15	0.17	0.31	0.70	0.70
Number of observations	90	89	86	81	80
Time span	1927–2016	1928–2016	1931–2015	1935–2016	1936–2016

Table 17: Regressions of the annual realized volatility of the S&P 500 and of the annual volatility of the 3-month Treasury bill on the posterior mean of the disaster intensity and several control factors. We generate a time series of the T-Bill yield volatility by considering at any point of time the 12 previous increments of the 3-month T-Bill yield and computing their empirical standard deviation; see Table 12 for summary statistics. The posterior disaster intensity is the sample path of $\mathbb{E}_t[\lambda_t]$ displayed in Figure 5. All data is sampled at an annual frequency. For control factors that are sampled monthly, we aggregate to annual time series by taking yearly averages. Given that our measures of volatility are backward-looking, we lag all the regressors by one unit of time. *** indicates significance at the 99.9% confidence level, * indicates significance at the 95% confidence level, and · indicates significance at the 90% confidence level.

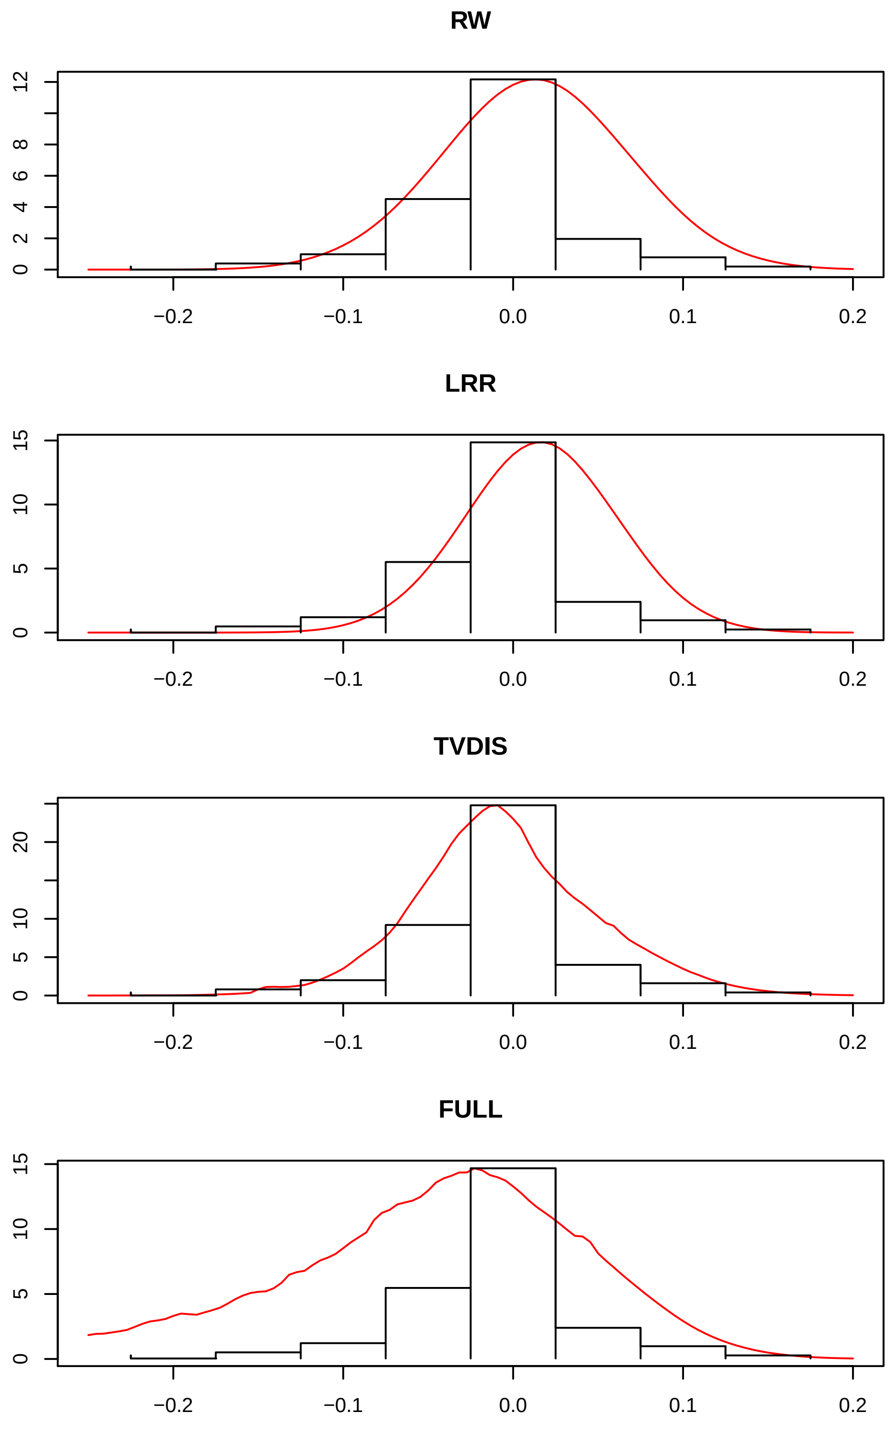

Figure 1: *Model-implied marginal distributions for Australia*. This figure shows in red the implied marginal distributions of consumption growth in Australia for different models. It also shows histograms of realized consumption growth in black. The marginal distributions are computed using the methodology of Guay and Schwenkler (2018) with the parameter estimates from Tables 4 and 5. We provide details in Appendix C.

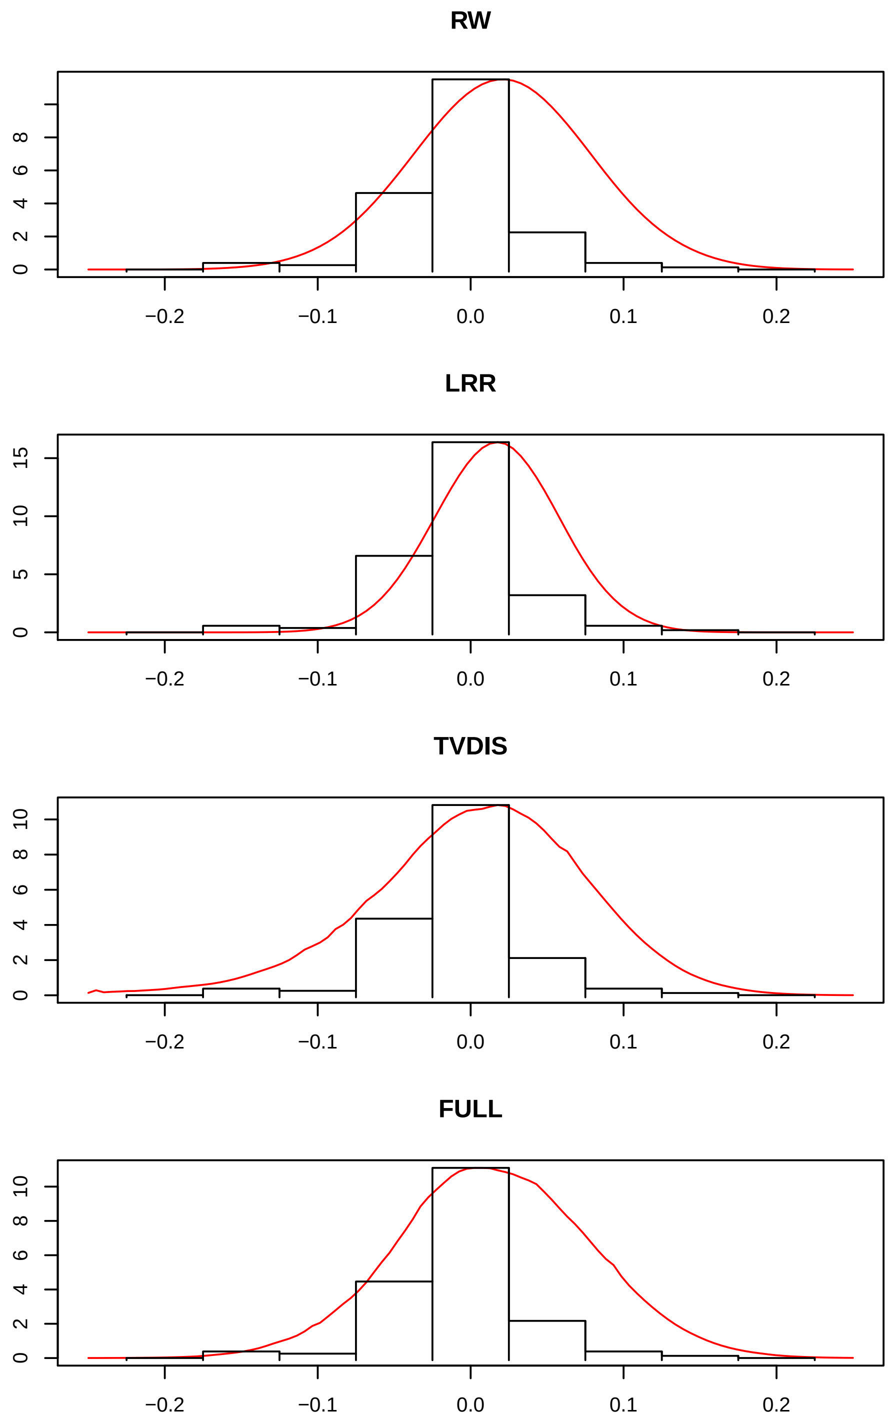

Figure 2: *Model-implied marginal distributions for Germany.* This figure shows in red the implied marginal distributions of consumption growth in Germany for different models. It also shows histograms of realized consumption growth in black. The marginal distributions are computed using the methodology of Guay and Schwenkler (2018) with the parameter estimates from Tables 4 and 6. We provide details in Appendix C.

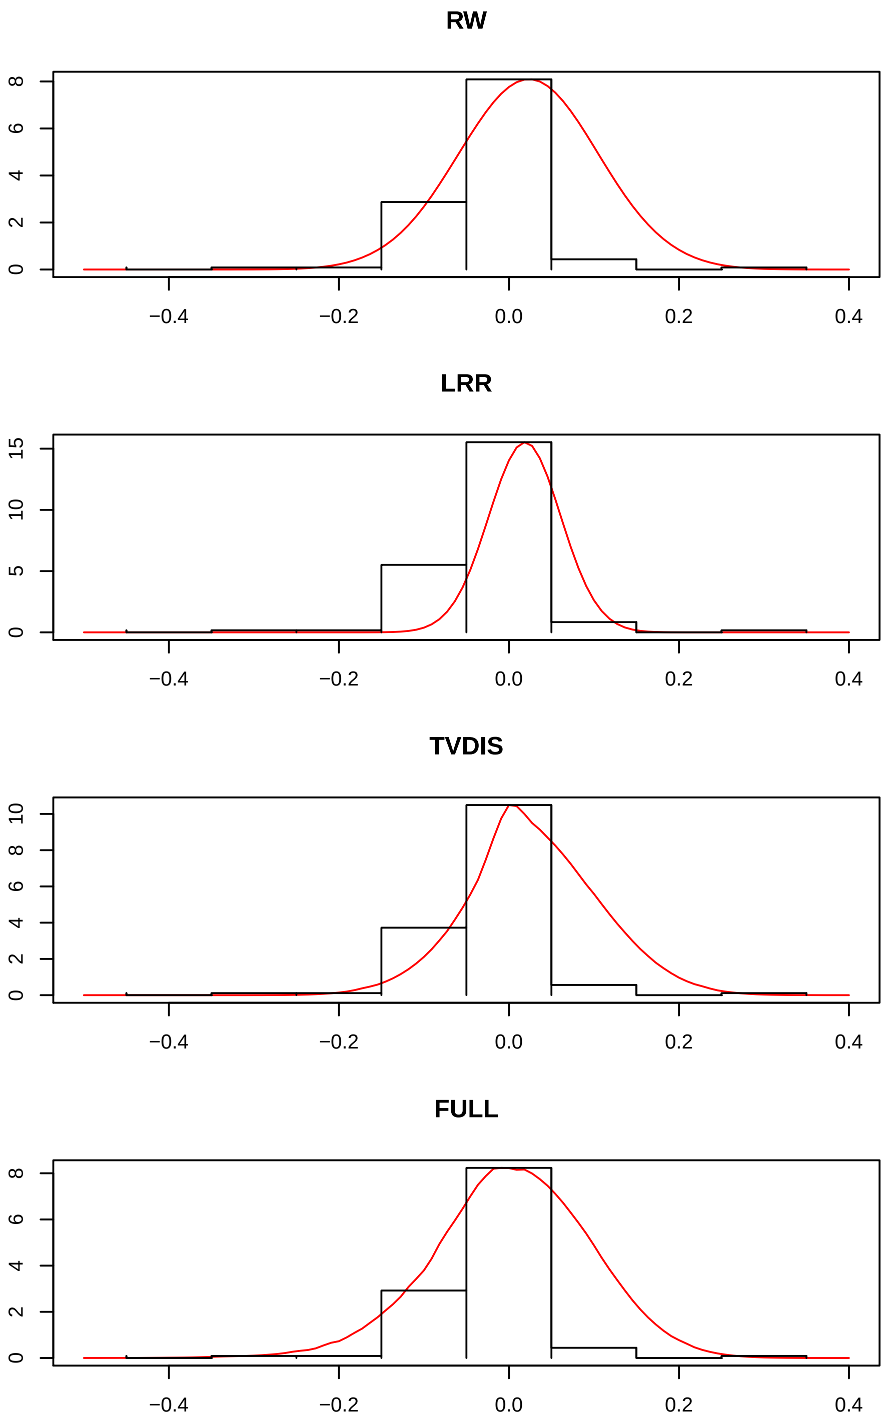

Figure 3: *Model-implied marginal distributions for Japan.* This figure shows in red the implied marginal distributions of consumption growth in Japan for different models. It also shows histograms of realized consumption growth in black. The marginal distributions are computed using the methodology of Guay and Schwenkler (2018) with the parameter estimates from Tables 4 and 7. We provide details in Appendix C.

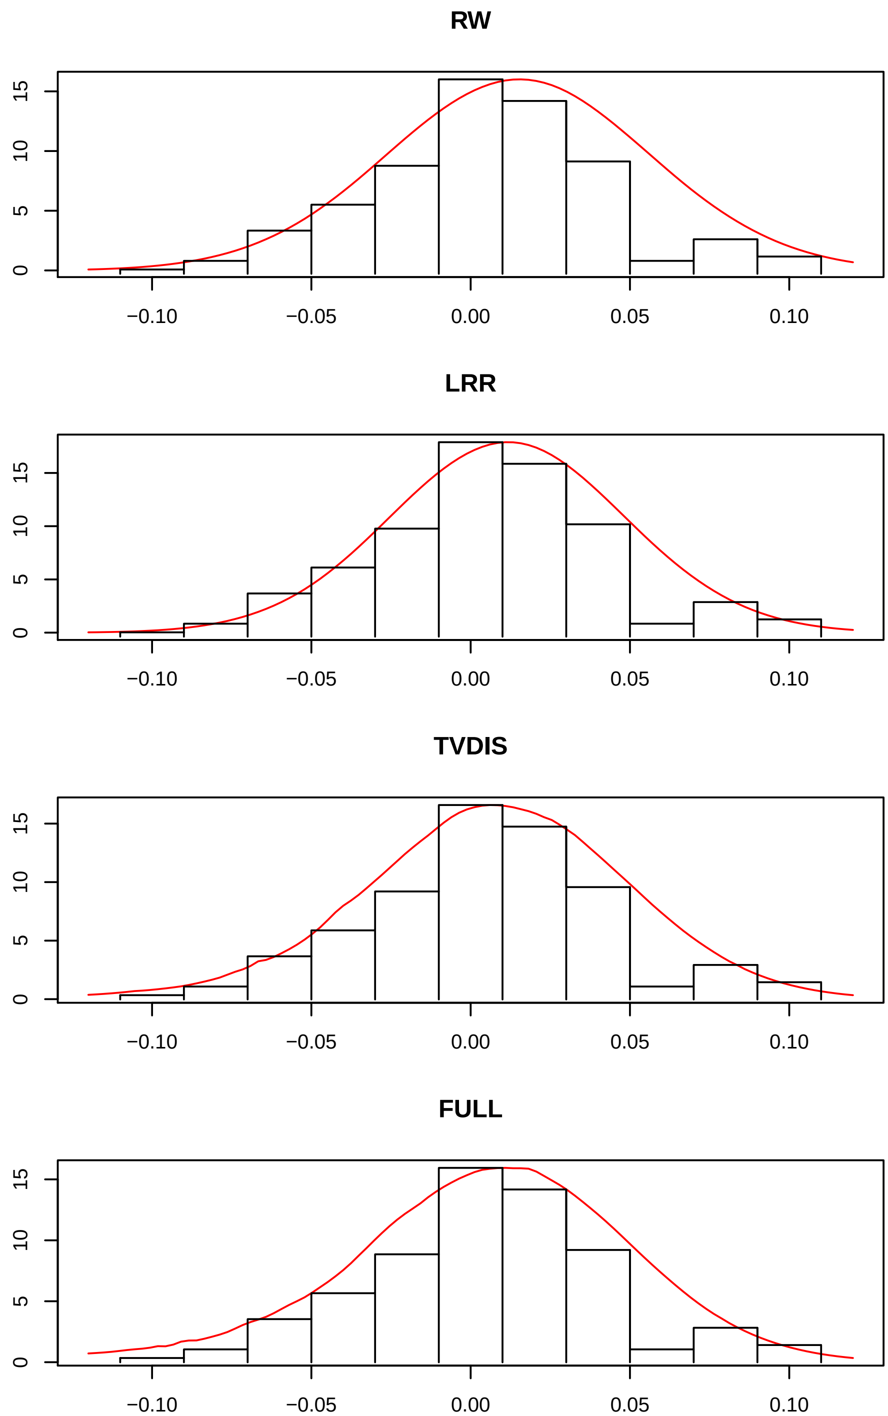

Figure 4: *Model-implied marginal distributions for the United States.* This figure shows in red the implied marginal distributions of consumption growth in the U.S. for different models. It also shows histograms of realized consumption growth in black. The marginal distributions are computed using the methodology of Guay and Schwenkler (2018) with the parameter estimates from Tables 4 and 8. We provide details in Appendix C.

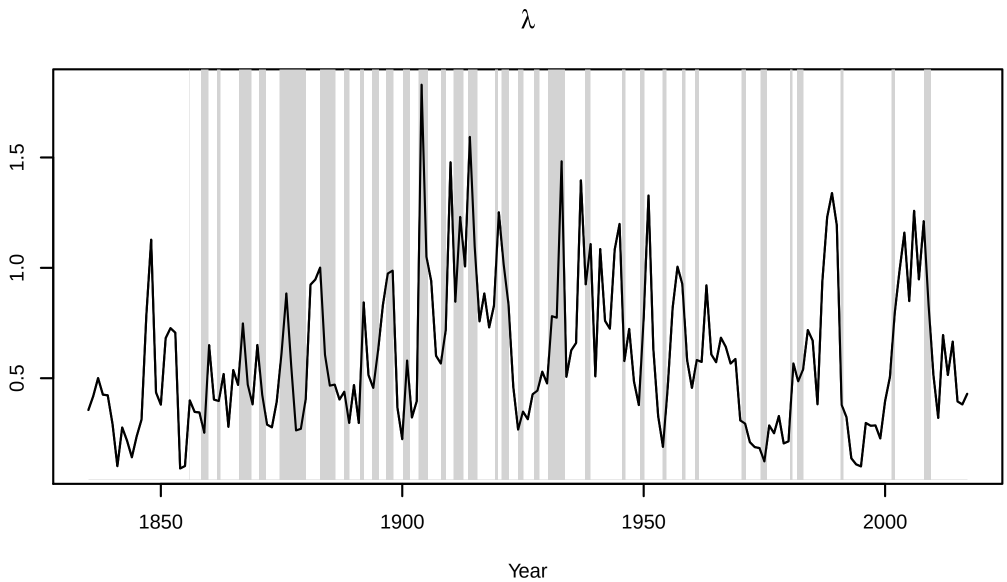

Figure 5: Posterior sample path of the disaster intensity ($\mathbb{E}_t[\lambda_t]$) for the United States (solid black line). We evaluate the posterior sample paths using the methodology of Guay and Schwenkler (2018) and the parameter estimates for the model "TVDIS" from Table 8. We also show NBER recessions as grey bars.

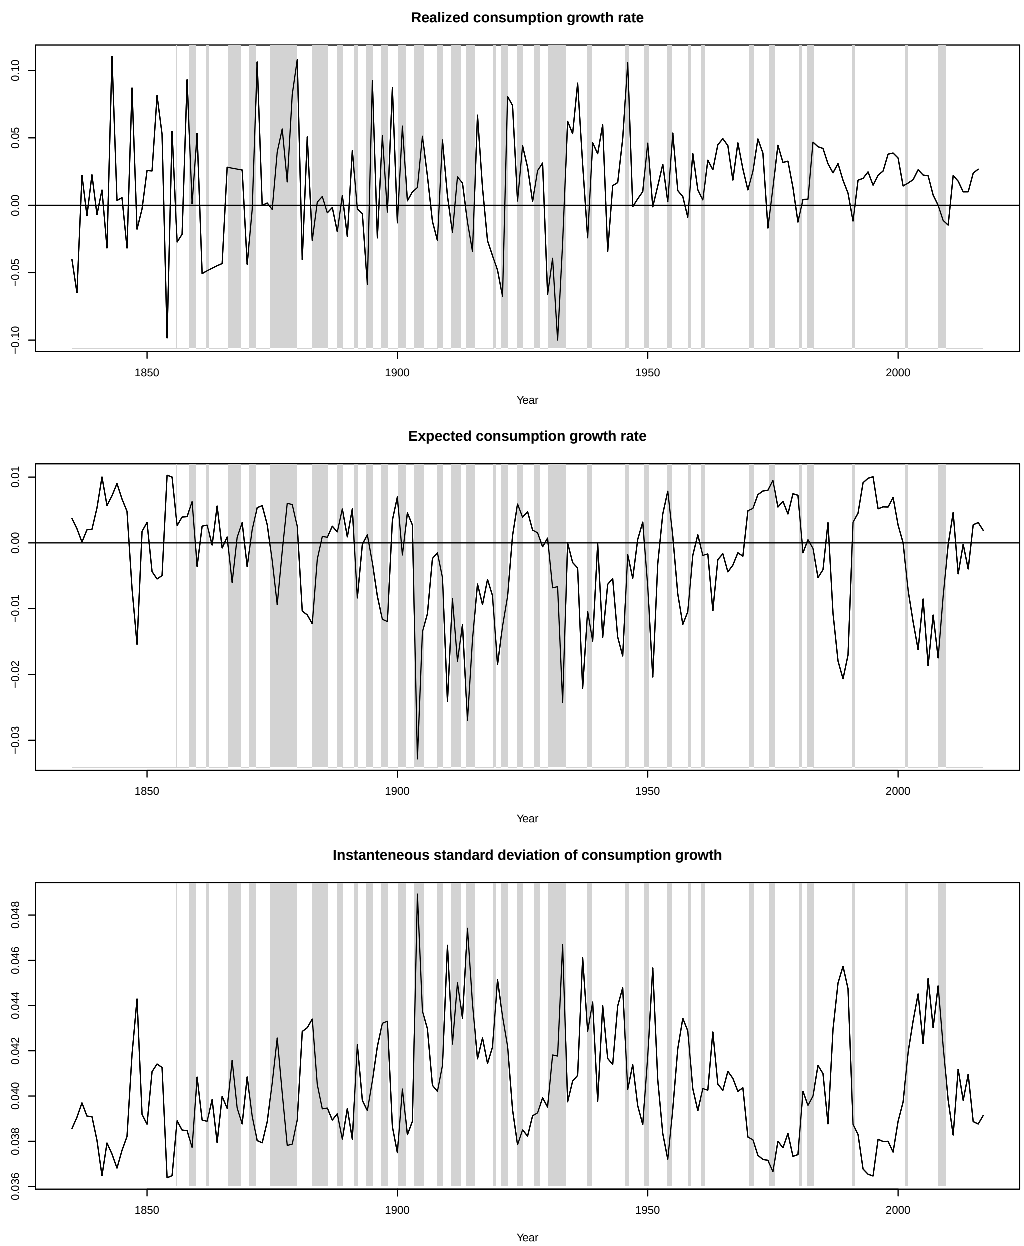

Figure 6: Posterior sample paths of the conditionally expected consumption growth rate $(\mu - \zeta \mathbb{E}_t[\lambda_t])$ and the conditional standard deviation of the instantaneous consumption growth rate $(\sqrt{\sigma^2 + \zeta^2 \mathbb{E}_t[\lambda_t]})$ for the United States. We use the posterior disaster intensity plotted in Figure 5 to compute these sample paths. We also show NBER recessions as grey bars. The top figure shows the realized annual consumption growth rates in the data.