Estimating the Dynamics of Consumption Growth

Gustavo Schwenkler*

March 13, 2018†

Abstract

We study the dynamics of consumption growth in a series of countries over a time span of 200 years. We seek to answer whether long-run risk or disasters are features of models that yield good fit to consumption data. To accomplish this goal, we develop a new methodology for the filtering and estimation of multivariate jump-diffusion processes in the presence of incomplete data and measurement errors. Our methodology is both statistically and computationally efficient, and enables the empirical analysis of previously intractable multidimensional models. Our estimates suggest that small and frequent disasters that arrive at a time-varying rate are a predominant feature of consumption data. Persistent stochastic volatility is also found to be a significant driver of consumption growth, especially in the United States.

*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, Olivier Scaillet, Viktor Todorov, the seminar participants at Boston University, the European Central Bank, Northwestern University, 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

Macro-finance models often postulate that consumption growth uncertainty is a significant driver of asset prices. A number of theories have been developed to explain anomalies in asset prices by adding novel features to standard models of consumption growth. For example, Bansal and Yaron (2004) posit that there exists a persistent latent component in consumption growth that drives variations in asset prices. Barro (2006) formulates a model with unexpected severe disasters to consumption growth that investors fear, again yielding variations in asset prices.

Many times, models of consumption growth are calibrated to match asset pricing moments. Such calibrations may be problematic if the target moments are not carefully chosen, and if the model of consumption growth is misspecified. The estimation of consumption growth models is challenging due to several layers of difficulty. First, consumption may depend in a complex way on latent factors – such as the long-run risk factors of Bansal and Yaron (2004) and the disaster factor of Barro (2006). The estimation of consumption growth models needs to address the fact that latent factors have to be filtered out based on the available data. Second, the joint distribution of consumption growth and the latent factors needs to be estimated jointly in a multidimensional model that allows for both small diffusive fluctuations and large disasters. The joint transition law of such multivariate jump-diffusion models is often intractable. Third, consumption data is observed only in discrete fashion, either on a monthly, quarterly, or annual basis. Consequently, the model of consumption growth need to be estimated from the available discrete data. Finally, consumption data may be contaminated with measurement errors. The question of how to carry out exact and efficient estimation for such complex models of consumption growth remains open.[1]

[1]The macro-finance literature recognizes this issue. 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

We answer this question by developing a novel methodology for the efficient estimation and filtering of multivariate jump-diffusion models in the presence of latent factors and measurement errors. Our econometric methodology makes it possible for us to estimate models of consumption growth that nest long-run risk and disaster specifications for a series of countries over a time span of 200 years. Our models feature latent persistent and transitory long-run risk factors that drive the expected growth rate and the conditional volatility of consumption as in the models of Bansal and Yaron (2004), Bansal et al. (2012), and Schorfheide et al. (2018). In addition, consumption growth may experience unexpected negative jumps as in the disaster models of Barro (2006), Barro and Jin (2011), and Barro and Ursúa (2017). The arrival rate of disaster may be time-variant, driven by an additional latent mean-reverting factor as in Wachter (2013). We estimate our models using annual data on real per capita consumer expenditures provided by Barro and Ursúa (2017) for Australia, Germany, Japan, and the United States. We also estimate several nested submodels. The simplest model posits consumption as a random walk, and does not allow for long-run risks or disasters. Intermediate models allow for transitory, persistent, or transitory and persistent long-run risk factors, but not for disasters. We also estimate models that allow for time-invariant or time-varying disasters, but do not account for long-run risks. All models are estimated both in the presence and the absence of measurement errors. By comparing the fit of the nested models to those of the complete model, we gain a complete picture of the relevance of the different features driving consumption growth.

The models we estimate are non-affine, and feature several latent factors as well as measurement errors. The estimation of such models using conventional methodologies is challenging. We circumvent these issues by extending the recently developed estimators of Guay and Schwenkler (2018) for multivariate jump-diffusion models. Guay and Schwenkler (2018) consider the problem of estimating the parameters of a multivariate jump-diffusion process based on discretely observed data in the presence of latent factors but without measurement errors. We extend the methodology of Guay and Schwenkler (2018) to allow for measurement errors, and show that the estimators we obtain for the model parameters

only recently turned."

and the measurement error variance are consistent, asymptotically normal, and asymptotically efficient. These results allow us to estimate all of our models using one and the same methodology, ensuring comparability across models. They also empower us with t-tests to assess the significance of individual model parameters. We further extend Guay and Schwenkler (2018) by developing likelihood ratio tests. Our likelihood ratio tests yield a quantitative approach to measure the incremental goodness-of-fit attained by each of the different model components.

Our parameter estimates suggest that time-varying disasters and stochastic volatility are significant drivers of consumption growth across the world, providing strong empirical support in favor of the models of Bansal and Yaron (2004) and Wachter (2013). Likelihood ratio tests corroborate these findings by showing that allowing for stochastic consumption volatility or time-varying disasters yields a significant improvement in the goodness of model fit for every country in our data sample. A model that allows both for stochastic volatility and time-varying disasters best fits the data, closely matching moments of international consumption data. Estimates of the posterior sample path of the disaster intensity and the stochastic volatility of consumption in the U.S. showcase that periods of negative consumption growth are associated with high consumption uncertainty and high disaster intensity. These empirical observations are consistent with the recent models of countercyclical consumption volatility, as in Andrei et al. (2016).

In spite of these positive results, our parameter estimates critically differ from estimates of the existing literature. Recently, Schorfheide et al. (2018) study a discrete-time model of consumption growth with long-run risks and measurement errors, but without disasters. The model of Schorfheide et al. (2018) is estimated via a Bayesian methodology using similar consumption data as ours. We pursue a frequentist approach, and estimate our model via maximum likelihood. Consistent with Schorfheide et al. (2018), we find that stochastic volatility is a significant driver of consumption growth in the United States. We also estimate shocks to stochastic consumption volatility to be similarly persistent as Schorfheide et al. (2018). Where we differ is in our estimates of the magnitudes of these shocks: We estimate the volatility of consumption growth volatility to be about 10 times smaller than Schorfheide et al. (2018).[2] We also cannot find the persistent component of expected consumption growth to be significantly large or time-varying. These finding lie at odds with estimates obtained by calibrating long-run risks models to asset pricing data, such as in Bansal and Yaron (2004), and suggest that calibrating a consumption growth model to asset pricing data may overstate the role of a persistent component in the expected consumption growth rate. Our findings on the insignificance of the persistent component of expected consumption growth are consistent with novel evidence by Beeler and Campbell (2012), who find only little evidence of predictability in U.S. consumption.

We find strong evidence in support of the occurrence of disasters across the world. However, the disasters we estimate are small and frequent: On average, we find that a disaster of about -3% occurs once every 13 months in one of the countries in our data. These estimates stand in stark contrast to the calibrated parameter values of Barro (2006), Barro and Jin (2011), and Barro and Ursúa (2017), who claim that disasters of about -20% occur one every 25 years. Our estimates imply similarly frequent and mild disasters as estimated by Backus et al. (2011), who look at U.S. equity option data to infer on the dynamics of consumption disasters. Going beyond Backus et al. (2011), we also allow for the rate of disaster arrival to vary over time in a mean-reverting fashion as in Wachter (2013). We find strong evidence that disaster rates rates vary strongly over time in all countries in our data. Shocks to the rate of disaster arrival are estimated to be similarly persistent as in Wachter (2013): A shock to the disaster rate in our model has a half-life of about 6.57 years, compared to 8.66 years in Wachter (2013). In contrast to Wachter (2013), however, our estimates imply smaller and more frequent disasters. Our results on the properties of consumption disasters are consistent with Julliard and Ghosh (2012), who non-parametrically estimate the frequency of disasters in the U.S. and several OECD countries, and conclude that this frequency is much higher than implied by the estimates in the disaster literature.

An analysis of the model-implied marginal distribution and moments of consumption growth shows that the main driver of the differences between our estimates and those of

[2] Our estimate of the volatility of consumption growth volatility is more in line with the estimate of Bansal et al. (2012).

Barro (2006), Barro and Jin (2011), Barro and Ursúa (2017), and Wachter (2013) is that we do not impose any left-tail moment match. As is standard in the disaster literature, the aforementioned studies calibrate the distribution and average frequency of disasters to match large declines in historical consumption data. We, on the other hand, estimate the distribution and frequency of disasters jointly with all other model parameters via maximum likelihood. Our analysis shows that a model that incorporates time-varying disasters but neglects long-run risks tends to overemphasize moments related to the left-tail of the empirical consumption growth distribution. In contrast, a model that incorporates long-run risks but neglects disasters tends to overemphasize moments related to the center of the consumption growth distribution. A full model that incorporates both time-varying disasters and long-run risks balances out these two forces, and yields estimates that are associated with small moment-matching errors. Our results have important implications for the modeling of consumption growth in asset pricing studies. They suggest that it is necessary to unify long-run risk and disaster models to properly capture the distribution of historical consumption data.

The rest of this paper is organized as follows. Section 2 introduces our models. Section 3 postulates our estimation methodology, and describes the data. Sections 4 presents and discusses our parameter estimates. A comparative analysis of the different nested models is carried out in Section 5.

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 large and 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 in order to be able to exploit recently developed tools 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, then the evaluation of the likelihood would be difficult. It would involve several approximations that may yield inefficient or biased parameter estimates.[3]

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

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

Here, W_1 is a standard Brownian motion, and

$$J_t = -\sum_{n \geq 1} \zeta E_n 1_{\{T_n \leq t\}} \tag{2}$$

is a pure-jump process with jump times $(T_n)_{n\geq 1}$ and stochastic intensity λ_t that we will specify below. The factors g_t, v_t, and λ_t introduce time variation in the conditional growth rate, volatility, and jump frequency of consumption. We specify these factors in the following sections.

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

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

$$v_t = \exp(vX_{3,t}) \tag{3}$$

for a process X_3 that solves the stochastic differential equation

$$\mathrm{d}X_{3,t} = -\kappa_3 X_{3,t} + \mathrm{d}W_{3,t}. \tag{4}$$

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.[4] The factor v_t scales the conditional volatility of consumption growth up and down over time around the baseline value of σ.

The parameter κ_3 gives an inverse measure of the persistence of the stochastic volatility factor: We could rewrite (4) as an AR(1) model, in which case $e^{-\kappa_3\Delta}$ would be the autoregressive coefficient. The parameter 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 growth. With this factor we model persistent long-run risk shocks as in Bansal and Yaron (2004). We assume that

$$g_t = \phi X_{2,t} \tag{5}$$

for a stochastic process $X_{2,t}$ that solves the stochastic differential equation

$$\mathrm{d}X_{2,t} = -\kappa_2 X_{2,t} + v_t \mathrm{d}W_{2,t} \tag{6}$$

for an independent standard Brownian motion W_2, an $X_{2,0} = 0$.[5] The persistent long-run risk factor is centered around zero but does not have Gaussian conditional distributions due to the presence of the stochastic volatility factor v_t in equation (6). Still, $\mathbb{E}[g_t] = 0$ for all $t \geq 0$. The parameter μ in (1) therefore measures the average consumption growth

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

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

rate, and ϕ measures the sensitivity of the expected consumption growth rate to long-run risk shocks. When $\phi = 0$, then there are no persistent long-run risk shocks to the expected growth rate of consumption. When $\phi > 0$, the parameter κ_2 gives an inverse measure of the persistence of a long-run growth risk shock. Equation (6) 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:

$$\Delta X_{1,T_n} = \Delta J_{T_n} = -\zeta E_n,$$

where 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).

Disasters in our model occur at the stochastic rate λ_t, which measures the conditional rate of disaster arrival (i.e., the conditionally expected number of disaster over a unit of time). We assume that

$$\lambda_t = \ell_0 \exp\left(\ell_1 X_{4,t}\right) \tag{7}$$

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

$$\mathrm{d}X_{4,t} = -\kappa_4 X_{4,t}\mathrm{d}t + \mathrm{d}W_{4,t}. \tag{8}$$

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

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

model of Wachter (2013). The arrival rate of disasters is time-invariant whenever $\ell_1 = 0$. The parameter κ_4 inversely measures the persistence of disaster intensity shocks: In an AR(1) formulation of (8), $e^{-\kappa_4 \Delta}$ would be the autoregressive coefficient. The parameter ℓ_1 captures how strongly the disaster intensity varies over time around the baseline disaster intensity ℓ_0. In contrast, the parameter ζ measures the average magnitude of a disaster. There are no disasters whenever $\zeta = 0$ and $\ell_0 = 0$.

3 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 consumption growth, long-run growth risk, stochastic volatility, and disasters. 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 1. We will use observations of log-consumption to estimate the parameters via maximum likelihood.

3.1 Data structure & measurement errors

We assume that the data is observed every Δ units on time, and that there are $m + 1$ data points available for inference. We do not assume that the available data are clean observations of the underlying log-consumption process. Instead, we make the assumption that the data is contaminated with potential measurement errors. At any observation time point $i\Delta$, the observed data is

$$\tilde{X}_{1,i\Delta} = X_{1,i\Delta} + \epsilon \nu_i, \tag{9}$$

where $(\nu_k)_{k \geq 0}$ is a sequence of i.i.d. standard normal noise variables. The parameter ϵ measures the standard deviation of measurement errors. When $\epsilon = 0$, there are no measurement errors. Controlling for measurement errors may be important for obtaining accurate parameter estimates with small standard errors.[7] Because of this, we will also estimate the parameter ϵ from the data.

[7] This is highlighted by Schorfheide et al. (2018), who emphasize that meaningful estimates of the persistence of long-run risk shocks can only be obtained when controlling for measurement errors.

In total, the data that is available for inference is

$$\mathsf{D}_m = \{\tilde{X}_{1,0}, \tilde{X}_{1,\Delta}, \ldots, \tilde{X}_{1,m\Delta}\}.$$

There is no data on the long-run growth risk factor g_t, the stochastic volatility factor v_t, the disaster intensity factor λ_t, and the error variables $(\nu_i)_{1\le i\le m}$. These processes are latent. The presence of latent factors poses significant challenges for the estimation of the model. We address these challenges in the following section.

3.2 Econometric methodology

Our goal is to estimate the joint dynamics of X_1, g, v, and λ, as well as the measurement error standard deviation ϵ. Due to our modeling choices in (1)-(9), the estimation problem is equivalent to estimating the dynamics of the multidimensional process $X = (X_1, X_2, X_3, X_4)$ with measurement errors. The process X is a multivariate jump-diffusion process that satisfies the following stochastic differential equation:

$$\mathrm{d}X_t = \begin{pmatrix} \mu + \phi X_{2,t} \\ -\kappa_2 X_{2,t} \\ -\kappa_3 X_{3,t} \\ -\kappa_4 X_{4,t} \end{pmatrix} \mathrm{d}t + \begin{pmatrix} \sigma e^{vX_{3,t}} & 0 & 0 & 0 \\ 0 & e^{vX_{3,t}} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \mathrm{d}W_t + \mathrm{d}J_t \tag{10}$$

for a four-dimensional standard Brownian motion $W = (W_1, W_2, W_3, W_4)$, and a non-explosive jump process

$$J_t = \sum_{n=1}^{N_t} \begin{pmatrix} -\zeta E_n \\ 0 \\ 0 \\ 0 \end{pmatrix} \tag{11}$$

with intensity $\lambda_t = \ell_0 e^{\ell_1 X_{4,t}}$. The process X specified in (10)-(11) is Markovian. If we had complete data that included observations $(X_{2,i\Delta}, X_{3,i\Delta}, X_{4,i\Delta})$ of the latent factors and the error variables ν_i for all $0 \le i \le m$, the likelihood at a given parameter θ would be the product of the transition density of X and the likelihood of the measurement errors:

$$\mathcal{L}_m^c(\theta) = \prod_{i=1}^{m} p_\Delta \left(\begin{pmatrix} \tilde{X}_{1,(i-1)\Delta} - \epsilon\nu_{i-1} \\ X_{2,(i-1)\Delta} \\ X_{3,(i-1)\Delta} \\ X_{4,(i-1)\Delta} \end{pmatrix}, \begin{pmatrix} \tilde{X}_{1,i\Delta} - \epsilon\nu_{i} \\ X_{2,i\Delta} \\ X_{3,i\Delta} \\ X_{4,i\Delta} \end{pmatrix}; \theta \right) \prod_{i=1}^{m} \phi\left(\nu_i; 0, 1\right). \quad (12)$$

Here, $\phi(\cdot; a, b)$ is the density of the normal distribution with mean a and standard deviation b, and p_Δ is the transition density of X specified by (10)-(11). Because the data does not include observations of the latent factors and the noise variables, the likelihood in our setting is the projection of the likelihood ratio onto the data space,

$$\mathcal{L}_m(\theta) = \mathbb{E}_{\theta^*}\left[\frac{\mathcal{L}_m^c(\theta)}{\mathcal{L}_m^c(\theta^*)} \,\middle|\, \mathsf{D}_m \right]. \quad (13)$$

Here, θ^* is the true data-generating parameter, and $\mathbb{E}_\theta$ denotes the expectation operator under the measure $\mathbb{P}$ when the parameter governing (10)-(11) is θ. Going forward, we will write $\mathbb{P}_\theta$ when referring to the measure $\mathbb{P}$ when the underlying parameter is θ.

The likelihood (13) is generally intractable. This is due to two reasons. First, the evaluation of the likelihood requires that we evaluate the transition density p_Δ of the multivariate jump-diffusion process specified in (10)-(11). The model (10)-(11) is not affine, which makes the evaluation of the density p_Δ difficult using standard tools.[8] Second, the likelihood (13) 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 processes as the one in (10)-(11).

We circumvent these difficulties by exploiting the recently developed methodology of Guay and Schwenkler (2018) for the efficient estimation of multivariate jump-diffusion processes. Inspired by Guay and Schwenkler (2018), we 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:

i) the sample paths $(X_{2,i\Delta})_{0\leq i\leq m}$, $(X_{3,i\Delta})_{0\leq i\leq m}$, and $(X_{4,i\Delta})_{0\leq i\leq m}$ of the latent factors are realizations of random walks that are independent of each other and of X_1,

[8] A similar model albeit in discrete time and without disasters is analyzed by Schorfheide et al. (2018).

ii) the realizations $(X_{1,i\Delta} - X_{1,(i-1)\Delta})_{0\le i\le m}$ of consumption growth are i.i.d. normally distributed with mean μ^* and standard deviation σ^*, and

iii) there are no measurement errors.

Such a measure $\mathbb{P}^*$ is equivalent to $\mathbb{P}_\theta$ for any parameter θ, and can be constructed via a standard Radon-Nikodym density.[9] Indeed, we can show that

$$\mathcal{L}_m(\theta) \propto \mathbb{E}^*\left[\frac{\mathcal{L}^c_m(\theta)}{\mathcal{L}^*_m}\,\middle|\, \mathsf{D}_m\right], \quad \text{where} \tag{14}$$

$$\mathcal{L}^*_m = \prod_{i=1}^m \phi\left(X_{2,i\Delta}; X_{2,(i-1)\Delta}, 1\right)\phi\left(X_{3,i\Delta}; X_{3,(i-1)\Delta}, 1\right)\phi\left(X_{4,i\Delta}; X_{4,(i-1)\Delta}, 1\right)$$
$$\times \prod_{i=1}^m \phi\left(\tilde{X}_{1,i\Delta} - \tilde{X}_{1,(i-1)\Delta}; \mu^*, \sigma^*\right)\prod_{i=1}^m \phi\left(\nu_i; 0, 1\right)$$

and $\mathbb{E}^*$ denotes the expectation operator under the new reference measure $\mathbb{P}^*$. The representation (14) facilitates the estimation of the likelihood via simulation. Whenever the density p_Δ is known, a simulation-based estimator of (14) can be constructed by generating i.i.d. samples of $(X_{2,i\Delta}, X_{3,i\Delta}, X_{4,i\Delta}, \nu_i)_{0\le i\le m}$ under the measure $\mathbb{P}^*$. This task can be easily accomplished because the latent factors follow random walk models under $\mathbb{P}^*$. Our approach for estimating the likelihood function extends the approach of Guay and Schwenkler (2018) by allowing for measurement errors.

We apply the methodology of Guay and Schwenkler (2018) to estimate the density p_Δ. Guay and Schwenkler (2018) derive a simulation-based estimator $\hat{p}_\Delta$ of the transition density p_Δ. This 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. A key feature of the density estimator of Guay and Schwenkler (2018) is that it is unbiased for all arguments of the density:

$$\mathbb{E}\left[P(v, w; \theta, \mathsf{R})\right] = p_\Delta(v, w; \theta).$$

[9] The measure $\mathbb{P}^*$ is equal to the measure $\mathbb{P}_\theta$ with a parameter vector θ for which $\mu = \mu^*$, $\sigma = \sigma^*$, and $\phi = v = \zeta = \kappa_2 = \kappa_3 = \kappa_4 = \epsilon = 0$.

The density estimator has finite variance and can be computed in finite time. Because of this, we can replace the true density p_Δ in (12) with the simulated density $\hat{p}_\Delta$, and still obtain an unbiased estimator of the likelihood in (14).

We summarize the steps we take to estimate the likelihood (14) in Algorithm 3.1. The correctness of Algorithm 3.1 follows from Guay and Schwenkler (2018). The number K of Monte Carlo samples used in Algorithm 3.1 controls the accuracy of the likelihood approximation. Larger values of K are associated with more precise estimates of the likelihood.

Algorithm 3.1. *Fix an integer K for the total number of Monte Carlo replications to be used to estimate the likelihood (14).*

(1) Generate a sequence $\left(\mathsf{R}^{(k)}\right)_{k=1,\dots,K}$ of independent replications of the random vector R as in Guay and Schwenkler (2018).

(2) Generate K independent sample paths $\left((X^{(k)}_{2,i\Delta})_{i\leq m}\right)_{k=1,\dots,K}$, $\left((X^{(k)}_{3,i\Delta})_{i\leq m}\right)_{k=1,\dots,K}$, and $\left((X^{(k)}_{4,i\Delta})_{i\leq m}\right)_{k=1,\dots,K}$ of independent random walks.

(3) For $k = 1,\dots,K$, compute:

$$\mathcal{L}^{c,(k)}_m(\theta) = \prod_{i=1}^m p_\Delta\left(\begin{pmatrix} \tilde{X}_{1,(i-1)\Delta} - \epsilon\nu^{(k)}_{i-1} \\ X^{(k)}_{2,(i-1)\Delta} \\ X^{(k)}_{3,(i-1)\Delta} \\ X^{(k)}_{4,(i-1)\Delta} \end{pmatrix}, \begin{pmatrix} \tilde{X}_{1,i\Delta} - \epsilon\nu^{(k)}_{i} \\ X^{(k)}_{2,i\Delta} \\ X^{(k)}_{3,i\Delta} \\ X^{(k)}_{4,i\Delta} \end{pmatrix}; \theta\right)$$

and

$$\mathcal{L}^{*,(k)}_m = \prod_{i=1}^m \phi\left(X^{(k)}_{2,i\Delta}; X^{(k)}_{2,(i-1)\Delta}, 1\right)\phi\left(X^{(k)}_{3,i\Delta}; X^{(k)}_{3,(i-1)\Delta}, 1\right)\phi\left(X^{(k)}_{4,i\Delta}; X^{(k)}_{4,(i-1)\Delta}, 1\right)$$
$$\times \prod_{i=1}^m \phi\left(\tilde{X}_{1,i\Delta} - \tilde{X}_{1,(i-1)\Delta}; \mu^*, \sigma^*\right),$$

An unbiased estimator of the likelihood (14) is given by

$$\hat{L}^K_m(\theta) = \frac{1}{K}\sum_{k=1}^K \frac{\mathcal{L}^{c,(k)}_m(\theta)}{\mathcal{L}^{*,(k)}_m}. \tag{15}$$

We compute approximate maximum likelihood estimators $\hat{\theta}_m^K$ by maximizing the likelihood of Algorithm 3.1:

$$\hat{\theta}_m^K \in \arg\max_{\theta \in \Theta} \hat{L}_m^K(\theta)$$

Here, Θ is a parameter space that contains all admissible parameters for our model. We assume that $\mu \in [-0.25, 0.25]$, $\sigma \in [10^{-8}, 1]$, $\phi \in [0, 1]$, $\kappa_2 \in [0, 3]$, $v \in [0, 1]$, $\kappa_3 \in [0, 3]$, $\zeta \in [0, 3]$, $\ell_0 \in [10^{-8}, 1]$, $\ell_1 \in [0, 3]$, $\kappa_4 \in [0, 3]$, $\epsilon \in [0, 1]$, and fix the parameter space Θ as the product of these intervals. Under mild assumptions, it is possible to show that $\hat{\theta}_m^K$ is a consistent estimator of θ^* if $K \to \infty$ as $m \to \infty$, and that $\sqrt{m}(\hat{\theta}_m^K - \theta^*)$ is normally distributed with the same asymptotic variance-covariance matrix as true maximum likelihood estimators if $\frac{m}{K} \to 0$ as both $m \to \infty$ and $K \to \infty$; see Guay and Schwenkler (2018). This yields our parameter statistically efficient, and implies that standard t-tests for statistical significance apply in our setting.

3.3 Nested models

In addition to Model (9)-(11), we will also estimate and analyze several models nested in our specification; see Table 2. If $\phi = v = \zeta = 0$, then the resulting model has no long-run risk and no disasters. Such a model posits log-consumption as a random walk, and we will estimate its free parameters μ and σ from the data as well. A second model we will estimate is a model that only allows for persistent long-run growth risk by setting $v = \zeta = 0$. A long-run growth risk model has four free parameters, namely μ, σ, ϕ, and κ_2. A third model we will estimate only allows for stochastic consumption growth volatility. Such a model restricts $\phi = \zeta = 0$, and has 4 free parameters: μ, σ, v, and κ_2. A long-run risk model in the spirit of Bansal and Yaron (2004) is obtained by restricting $\zeta = 0$. Such a model has 6 free parameters: μ, σ, ϕ, κ_2, v, and κ_3. We will also estimate two types of disaster models. The simpler disaster model assumes that the disaster intensity is constant as in Barro (2006). Such model restricts $v = \phi = \ell_1 = 0$ 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 $\phi = v = 0$. We will estimate and analyze these nested models both in the absence

($\epsilon = 0$) and the presence ($\epsilon > 0$) of measurement error.

In addition to using t-tests to assess the significance of each model parameter, we will also use likelihood ratio tests to determine the incremental fit obtained by relaxing the restrictions of the nested models. Suppose that $\hat{\theta}^K_{m,R}$ and $\hat{\theta}^K_{m,U}$ are the parameter estimates computed with our methodology for a restricted model R and and unrestricted model U, respectively. The unrestricted model achieves a better fit to the data than the restricted model because it has more free parameter 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. We can show that the likelihood ratio statistic implied by our methodology,

$$\mathrm{LR}(U,R) = 2\left(\hat{L}^K_m\left(\hat{\theta}^K_{m,U}\right) - \hat{L}^K_m\left(\hat{\theta}^K_{m,R}\right)\right), \tag{16}$$

has a chi-squared asymptotic distribution as $m, K \to \infty$ and $\frac{m}{K} \to 0$. 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 tests next to the t-tests of Guay and Schwenkler (2018) empowers us with tools to carry out formal statistical analysis of the full model and the nested submodels.

3.4 The data

We estimate the models of Table 2 using consumption growth data. We use the Barro-Ursúa data set for this.[10] The data contain annual observations of real per capita consumer expenditure for a cross section of countries. We select the following countries for our analysis:

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

- 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.

4 Full model

Table 3 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 are significant across countries. We also see that the disaster intensity is estimated to be significantly time-varying in all countries. The disaster magnitude estimated by the full model is significantly large in Australia, Germany, and the U.S.. The disaster magnitude estimated by the full model with measurement errors is significantly large in Germany and Japan. Stochastic consumption volatility is significant in the U.S., but not in the other countries we consider in our data. We cannot find any evidence of a significant persistent component in consumption growth ($\phi > 0$). We can also not identify significantly large measurement errors ($\epsilon > 0$).

Figure 1 shows the sample path of realized consumption growth in the U.S., and compares this path to the posterior sample path of the persistent consumption growth component ϕg_t, the stochastic volatility σv_t of consumption growth, and the disaster intensity λ_t. We see that the earlier part of our sample period was associated with high consumption volatility, and that the posterior mean of consumption volatility has come down in the past 50 years. Similarly, the posterior mean of the disaster intensity was high in the earlier parts of our sample periods, and has stabilized around a value of 1.3 disasters per year in the last 120 years. We also see that the posterior mean of consumption volatility and the posterior mean of the disaster intensity spike up during periods of negative consumption growth, suggesting that such periods are associated with high economic uncertainty and high economic risks. One of the periods with the largest negative consumption growth rates is associated with the great depression. During that period we also observe the highest posterior mean of consumption growth volatility. These observations are consistent with the learning model of Andrei et al. (2016). In that model, agents infer from bad productivity shocks that productivity will remain low for a long stretch of time, resulting in more volatile consumption. This yields countercyclical consumption volatility patterns as those reflected in the posterior means of Figure 1.

4.1 Relation to existing estimates in the literature

Our estimated suggest that disasters are a significant driver of consumption growth in all four countries in our data sample. There is some heterogeneity in the disaster estimates across countries: Some countries experience larger disasters, other countries experience disasters more frequently. For example, our estimates suggest that an average disaster of -1.20% arrives every 1.12 years in Japan, while in Australia an average disaster of -4.66% arrives every 7 months. Across the board, however, disaster magnitudes are much smaller and more frequent than calibrated by Barro (2006), Barro and Jin (2011), Barro and Ursúa (2017), and others. For example, using the same data as we use, Barro and Jin (2011) estimate that a disaster of median size of -19% arrives once every 26 years in the U.S.. In contrast, we estimate that in the U.S. a disaster of median size of -0.01% arrives once every 1.27 years.[11] Our estimates of the average disaster magnitude and the average disaster frequency for the U.S. are comparable to the estimates of Backus et al. (2011). Using options data, Backus et al. (2011) estimate than an average U.S. consumption disaster of -0.74% arrives once every 9 months.

We also find that the arrival rate of consumption disasters is strongly time-varying. This is illustrated in Figure 1, which shows that the posterior mean of the U.S. disaster intensity can spike up during periods of low consumption growth. Our estimates

[11] The long-run mean of the time-varying intensity in our model is $\ell_0 \exp(\frac{1}{2}\frac{\sigma^2}{2\kappa_4})$ due to the Gaussian distribution of X_4.

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, 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 0.11, and the corresponding estimate of Wachter (2013) is 0.08. In contrast to Wachter (2013), our estimate of the average U.S. disaster intensity is much larger. Wachter (2013) estimates that a disaster occurs in the U.S. on average every 28 years, similarly as in Barro (2006), Barro and Jin (2011), Barro and Ursúa (2017), and others. We find that a disaster occurs every 1.27 years, which is about as frequent as in Backus et al. (2011).

Many of the differences between our disaster estimates and those established by the literature can be explained by differences in the estimation approaches. In Barro (2006), Barro and Jin (2011), Barro and Ursúa (2017), and Wachter (2013), disasters are calibrated to match the left tail of consumption growth. In contrast, we estimate the magnitude and frequency of disasters together with the other model parameters via maximum likelihood. By doing so, we do not restrict our econometric methodology to explicitly fit certain moments of the consumption growth distribution. Instead, we allow the maximum likelihood methodology to estimate the model parameters as to minimize the distance between the model-implied and the empirical distribution of the data. As showcased in Table 4, the moments implied by our parameter estimates closely reflect moments in the data. We conclude that by enforcing a match of left-tail moments of consumption growth, the resulting disaster severity and frequency may be overly conservative. This conclusion is consistent with Julliard and Ghosh (2012), who non-parametrically estimate the probability of experiencing a disaster on any given year in the U.S. and a series of OECD countries without imposing a left-tail match, and find that this probability is much larger than implied by Barro (2006).

Schorfheide et al. (2018) estimate a generalization of the long-run risk model of Bansal and Yaron (2004) via a Bayesian methodology using only consumption data for the United States. Bansal et al. (2012) calibrate a long-run risk model to annual U.S. consumption and asset pricing data. Both of these studies find supporting evidence in favor of a persistent consumption growth component as well as of persistent stochastic consumption volatility. In contrast to the aforementioned studies, we only find empirical support of persistent stochastic consumption growth volatility. If we were to translate our estimate of κ_2 to a monthly AR(1) coefficient analogous to the one estimated by Bansal et al. (2012) and Schorfheide et al. (2018) for stochastic consumption growth volatility in the U.S., we would obtain an estimate of 0.999. This estimates is the same as the analogous estimate of 0.999 by Bansal et al. (2012), and falls inside of the 90% credible band estimate by Schorfheide et al. (2018) that reaches from 0.976 to 0.999. The estimates of Schorfheide et al. (2018) imply 90% confidence bands for our parameter v ranging between 0.092 and 0.190. Our estimate of v for the U.S. of 0.011 lies below this credible band, and is slightly higher than the analogous estimate of 0.006 of Bansal et al. (2012). The differences between our estimates and the estimates of Bansal et al. (2012) and Schorfheide et al. (2018) may be explained by the fact that we also control for disasters when evaluating the significance of stochastic volatility for modeling consumption growth.

We cannot find the parameter ϕ to be significantly positive in any country, suggesting that there is only insignificant time-variation in expected consumption growth rates. This finding lies in contrast to estimates of ϕ calibrated from asset pricing data in Bansal and Yaron (2004), Bansal et al. (2012), and others. In those studies, the parameter ϕ plays a key role for matching asset pricing moments. Our findings suggest that consumption growth data may not support strongly time-varying expected consumption growth rates, and that the role of a persistent component in the expected consumption growth rate may be overstated if calibrated from asset pricing data. These conclusions are consistent with evidence of Beeler and Campbell (2012), who find only little evidence of predictability in U.S. consumption growth data.

We also cannot find any evidence of significant measurement errors in any country. This lies in agreement with Schorfheide et al. (2018), who can only identify minor measurement errors for annual consumption growth in the U.S. using a Bayesian estimation approach. Similarly as in Schorfheide et al. (2018), we find that controlling for measurement errors can help with the identification of some model parameters. For example, in the U.S. we estimate much smaller and less significant disasters in a model that controls for measurement errors than in a model that does not control for measurement errors. In Japan, we can establish the significance of the average disaster magnitude parameter only when controlling for measurement errors. These observations suggest that measurement errors may be small in annual consumption data, but that controlling for the influence of measurement errors is an important econometric tool to accomplish good identification.

5 Nested models

Tables 5 through 8 report the parameter estimates for the nested models fitted to Australian, German, Japanese, and U.S. consumption data. These estimates provide several new insights. On the measurement error front, we still find little evidence of significant errors; only in Japan do we sometimes find the parameter ϵ to be significantly large. The models that allow for time-invariant disasters (Models "DIS" and "DIS+ME") cannot identify significant disasters in any country. This observation lies at odd with the disaster estimates of Barro (2006), Barro and Jin (2011), and Barro and Ursúa (2017), and provides further evidence in support of the hypothesis that calibrating disasters to the left tail of the empirical consumption growth distribution may overstate the role of disasters in historical consumption data.

We identify significant disasters in Australia and Germany only when allowing for time-variation in the disaster intensity (Models "TVDIS" and "TVDIS+ME"). 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 3. We cannot find the average disaster magnitude to be significantly large in Japan and the U.S., even though the baseline disaster rate, the disaster intensity volatility, and the disaster intensity speed of reversion are found to be significantly large for these countries. The disaster estimates for Japan and the U.S. in the restricted models that allow for time-varying disaster intensities ("TVDIS" and "TVDIS+ME") are also similar to the estimates reported for the full models in Table 3. 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).

In Australia, Germany, and Japan, the nested models that allow for stochastic volatility (Models “LRVR”, “LRVR+ME”, “LRR”, and “LRR+ME”) estimate significant time-variation and persistence in the volatility of consumption growth. The estimates of the speed of reversion κ_3 of stochastic volatility in these countries imply monthly AR(1) coefficients that lie inside of the 90% credible bands estimated for this parameter by Schorfheide et al. (2018) based on U.S. data. The models that allow for a persistent component in the expected consumption growth rate (Models “LRGR”, “LRGR+ME”, “LRR”, and “LRR+ME”) only find this component to be significantly large in Japan. For that country, our estimates of the speed of reversion κ_2 of the persistent consumption growth component imply a monthly AR(1) coefficient of 0.998, which is comparable to the monthly AR(1) coefficient estimate of Bansal et al. (2012) for U.S. data. The 90% credible band implied by Schorfheide et al. (2018) for their AR(1) coefficient ranges 0.923 to 0.978, suggesting that our estimated persistent consumption growth factor for Japan is more persistent than the analogous factor estimated by Schorfheide et al. (2018) for the United States. When we compare the estimates of the long-run risk parameters from the nested models in Tables 5 through 8 to the estimates of the full models in Table 3, we see that the significance of the long-run risk parameters disappears when we also control for time-varying disasters. This observation suggests that the role of long-run risks for explaining historical consumption data may be overstated in a model that neglects the influence of time-varying disasters.

5.1 Implied moments and distributions

The estimates of the nested models in Tables 5 through 8 suggest that models of consumption growth that do not control for all possible sources of risk may overstate certain model components. To assess the impact of such overstatements, in Tables 9 through 12 we report model-implied moments for the nested models and compare to 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, the variance, and the kurtosis. In contrast, the models that control for disasters but not for long-run risks match the skewness as well as the left tail of the empirical consumption growth distribution. The full model balances out these two forces, and generally achieves small moment-matching errors.

We reach similar conclusions when looking at the model-implied marginal distributions of consumption growth displayed in Figures 2 and 3 for Japan and the United States. The model that only includes time-varying disasters and neglects long-run risks tends to overstate the left tail of the empirical consumption growth distribution. The model that only includes long-run risks and neglects disasters tends to overstate the center of the empirical distribution. A full model that incorporates both long-run risks and time-varying disasters tends to perform best in matching the empirical distribution of U.S. consumption growth.

The model-implied moments and distributions provide a quantitative justification for the differences between our parameter estimates and those established by the previous literature. As modeling features, long-run risks and disasters target different parts of the consumption growth distribution. When one estimates a model that includes one feature but not the other, necessarily the active feature needs to be overstated in order to compensate for the missing feature. Our results suggest that models that estimate long-run risks in the absence of disasters will tend to overstate the role of long-run risks. Similarly, models that estimate disasters in the absence of long-run risks will overstate the role of disasters. A full model that incorporates time-varying disasters and long-run risks is necessary to achieve a good fit to historical consumption data. In this regard, our results provide validation of a claim by Ludvigson (2013) that states that "the allowance for disasters in standard models of risk provides an example of how superior specifications may potentially be obtained by combining elements of several consumption-based models."

5.2 Likelihood ratio tests

We finalize this section by measuring the incremental goodness-of-fit achieved by allowing for disaster and long-run growth risk in a consumption growth model. This will allow us to determine the relative importance of long-run risks and disasters for explaining historical consumption data. Table 13 reports the likelihood ratio statistics of (16). The results show that models that include only long-run risks or only time-varying disasters are preferred over simple random walk models for all countries in our data. In terms of the different components of long-run risks models – that is, persistent expected consumption growth and stochastic volatility – we find that the inclusion of stochastic volatility generally leads to a more significant increase in the goodness-of-fit of a model than the inclusion of a persistent long-run consumption growth component. However, we find that the inclusion of time-varying disasters leads to a higher and more significant increase in the goodness-of-fit of a model than the inclusion of both stochastic volatility and persistent consumption growth. The likelihood ratio tests do not provide support in favor of a time-invariant disaster model. We conclude that the better fit provided by a time-varying disaster model is primarily achieved through the introduction of time variation in the disaster rate rather than the introduction of rare disasters per se.

The likelihood ratio tests also show that introducing time-varying disasters to a model that includes long-run risks always leads to a significant increase in the goodness-of-fit of the model. In contrast, the inclusion of long-run risks to a model with time-varying disasters only leads to a significant increase in the fit to U.S. consumption data. The estimates of the full model reported in Table 3 indicate that this is primarily due to the inclusion of stochastic volatility rather than the inclusion of a persistent consumption growth component. This observation is consistent with our interpretation of the likelihood ratio tests for the models "LRGR+ME" and "LRVR+ME" which only allow for either persistent consumption growth or stochastic consumption volatility.

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. Stochastic volatility is also supported by the data, although to a lesser degree than time-varying disasters. Our results favor models of consumption growth that allow for mean-reverting and time-varying disaster intensities, such as in the model of Wachter (2013). They also provide some support for models that allow for persistent economic uncertain shocks in the form of stochastic volatility, such as in the model of Bansal and Yaron (2004).

References

Andrei, Daniel, Steffen Hitzemann and Alexandre Jeanneret (2016), Learning and consumption dynamics. Working Paper.

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 (2012), 'An empirical evaluation of the long-run risks model for asset prices', *Critical Finance Review* **1**(1), 183–221.

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.

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

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

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* . Forthcoming.

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

Julliard, Christian and Anisha Ghosh (2012), 'Can rare events explain the equity premium puzzle?', *The Review of Financial Studies* **25**(10), 3037–3076.

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.

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), $\sqrt{12}\sigma$ in Schorfheide et al. (2018), and μ in Wachter (2013).
σ	Annualized baseline consumption growth volatility in the absence of stochastic volatility	This parameter is the same as σ in Barro (2006) and Wachter (2013), and $\sqrt{12}\sigma$ in Bansal and Yaron (2004) and Schorfheide et al. (2018).
ϕ	Sensitivity of consumption growth to persistent long-run growth 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 long-run growth risk component	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 consumption growth volatility	This parameter is analogous to $\sqrt{12}\sigma_w$ in Bansal and Yaron (2004).
κ_3	Speed of reversion of stochastic consumption growth volatility	This parameter can be translated to a monthly AR(1) coefficient analogous to the parameter ν_1 in the long-run risk model of Bansal and Yaron (2004) or ρ_{h_c} in Schorfheide et al. (2018) by taking $e^{-\kappa_3/12}$.
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).
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).
ϵ	Standard deviation of measurement errors	This parameter is equivalent to the parameter σ_ϵ^a in Schorfheide et al. (2018).

Table 1: *Parameters*. This table summarizes the roles played by each of the parameters in Model (1)-(7).

Model	Interpretation	Free parameters	Restricted parameters
RW	Random walk model	(μ, σ)	$\phi = v = \zeta = \ell_1 = \epsilon = 0$
RW + ME	RW model with measurement errors	(ϵ, μ, σ)	$\phi = v = \zeta = \ell_1 = 0$
LRGR	Model that allows for a persistent long-run risk component driving the conditional growth rate of consumption	$(\mu, \sigma, \phi, \kappa_2)$	$v = \zeta = \ell_1 = \epsilon = 0$
LRGR + ME	LRGR model with measurement errors	$(\epsilon, \mu, \sigma, \phi, \kappa_2)$	$v = \zeta = \ell_1 = 0$
LRVR	Model that allows for a transitory long-run risk component driving the stochastic volatility of consumption growth	$(\mu, \sigma, v, \kappa_3)$	$\phi = \zeta = \ell_1 = \epsilon = 0$
LRVR + ME	LRVR model with measurement errors	$(\epsilon, \mu, \sigma, v, \kappa_3)$	$\phi = \zeta = \ell_1 = 0$
LRR	Model that allows for transitory and persistent long-run risk components in the conditional growth rate and the conditional volatility of consumption growth	$(\mu, \sigma, \phi, \kappa_2, v, \kappa_3)$	$\zeta = \ell_1 = \epsilon = 0$
LRR + ME	LRR model with measurement errors	$(\epsilon, \mu, \sigma, \phi, \kappa_2, v, \kappa_3)$	$\zeta = \ell_1 = 0$
DIS	Model that allows for time-invariant disaster risk	$(\mu, \sigma, \zeta, \ell_0)$	$\phi = v = \ell_1 = \epsilon = 0$
DIS + ME	DIS model with measurement errors	$(\epsilon, \mu, \sigma, \zeta, \ell_0)$	$\phi = v = \ell_1 = 0$
TVDIS	Model that allow for time-varying disaster risk	$(\mu, \sigma, \zeta, \ell_0, \ell_1, \kappa_4)$	$\phi = v = \epsilon = 0$
TVDIS + ME	TVDIS model with measurement errors	$(\epsilon, \mu, \sigma, \zeta, \ell_0, \ell_1, \kappa_4)$	$\phi = v = 0$
FULL	Complete model that allows for persistent and transitory long-run risk components, as well as time-varying disaster risk	$(\mu, \sigma, \phi, \kappa_2, v, \kappa_3, \zeta, \ell_0, \ell_1, \kappa_4)$	$\epsilon = 0$
FULL + ME	FULL model with measurement errors	$(\epsilon, \mu, \sigma, \phi, \kappa_2, v, \kappa_3, \zeta, \ell_0, \ell_1, \kappa_4)$	

Table 2: *Models*. This table summarizes the different models that we estimate in this paper.

	AU		DE		JP		US	
	FULL	FULL+ME	FULL	FULL+ME	FULL	FULL+ME	FULL	FULL+ME
μ	** 0.0194	** 0.0194	*** 0.0245	*** 0.0245	*** 0.0307	*** 0.0316	*** 0.0205	*** 0.0204
	(0.0066)	(0.0066)	(0.0057)	(0.0069)	(0.0035)	(0.0039)	(0.0021)	(0.0029)
σ	*** 0.0525	*** 0.0525	*** 0.0572	*** 0.0572	*** 0.0755	*** 0.0755	*** 0.0412	*** 0.0412
	(0.0048)	(0.0050)	(0.0036)	(0.0037)	(0.0038)	(0.0054)	(0.0015)	(0.0027)
ϕ	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0009	0.0009
	(0.0022)	(0.0022)	(0.0015)	(0.0016)	(0.0012)	(0.0012)	(0.0006)	(0.0015)
κ_2	· 0.0477	· 0.0477	* 0.0470	* 0.0471	** 0.0622	** 0.0592	*** 0.0557	*** 0.0560
	(0.0272)	(0.0272)	(0.0250)	(0.0250)	(0.0221)	(0.0219)	(0.0099)	(0.0111)
v	0.0008	0.0008	0.0036	0.0036	0.0000	0.0001	*** 0.0114	* 0.0114
	(0.0088)	(0.0088)	(0.0057)	(0.0058)	(0.0047)	(0.0068)	(0.0026)	(0.0050)
κ_3	0.0123	0.0123	0.0082	0.0081	0.0036	0.0021	** 0.0128	· 0.0128
	(0.0107)	(0.0107)	(0.0080)	(0.0080)	(0.0084)	(0.0084)	(0.0047)	(0.0079)
ζ	* 0.0466	* 0.0466	*** 0.0157	** 0.0388	0.0393	*** 0.0120	* 0.0574	0.0001
	(0.0209)	(0.0210)	(0.0036)	(0.0140)	(0.0484)	(0.0016)	(0.0280)	(0.0079)
ℓ_0	*** 1.7759	*** 1.7759	*** 0.7593	*** 0.7596	*** 0.8865	*** 0.8916	*** 0.7582	*** 0.7591
	(0.2687)	(0.2687)	(0.1134)	(0.0987)	(0.1658)	(0.1824)	(0.0646)	(0.0773)
ℓ_1	*** 0.5637	*** 0.5637	*** 0.4487	*** 0.4489	*** 0.4217	*** 0.4213	*** 0.4655	*** 0.4653
	(0.1028)	(0.1072)	(0.0456)	(0.0383)	(0.0561)	(0.0607)	(0.0265)	(0.0418)
κ_4	* 0.0867	* 0.0867	** 0.0918	** 0.0918	** 0.0844	** 0.0852	*** 0.1060	*** 0.1055
	(0.0395)	(0.0397)	(0.0342)	(0.0348)	(0.0286)	(0.0301)	(0.0200)	(0.0161)
ϵ		0.0000		0.0000		0.0000		0.0002
		(0.0043)		(0.0031)		(0.0041)		(0.0021)
Likelihood	352.72	352.72	487.00	487.00	372.89	372.93	631.76	631.76

Table 3: 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 $\hat{L}_m^K(\theta)$ of Algorithm 3.1 with $K = 10^3$. We use the Nelder-Mead algorithm in R to compute the maximizer of the simulated likelihood. Given in parenthesis are asymptotic standard errors computed as outlined by Guay and Schwenkler (2018). 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.

	AU		DE		JP		US	
	Data	FULL+ME	Data	FULL+ME	Data	FULL+ME	Data	FULL+ME
Mean	0.014	0.010	0.017	−0.004	0.022	0.022	0.015	0.018
Standard deviation	0.050	0.075	0.053	0.071	0.069	0.062	0.038	0.045
Skewness	−0.982	−0.278	−0.550	−0.126	−1.509	−0.015	−0.071	−0.117
Kurtosis	7.910	1.726	7.264	2.865	19.887	2.589	3.528	3.154
Probability of a decline of more than 5%	0.074	0.268	0.057	0.249	0.059	0.137	0.038	0.063
Probability of a decline of more than 10%	0.028	0.079	0.038	0.027	0.022	0.023	0.005	0.009
First-order autocovariance	0.000	0.001	0.001	0.004	0.005	0.001	0.000	0.001

Table 4: Model-implied moments of consumption growth for the different countries in our analysis, for both the full model ("FULL") and the full model with measurement errors ("FULL+ME"). We compute model-implied moments using the methodology of Guay and Schwenkler (2018) with $K = 10^5$. We compare these model-implied moments to the corresponding moments in the data. All moments are measured in annual terms.

	AU											
	RW	RW+ME	LRGR	LRGR+ME	LRVR	LRVR+ME	LRR	LRR+ME	DIS	DIS+ME	TVDIS	TVDIS+ME
μ	*** 0.013	*** 0.013	*** 0.017	*** 0.017	*** 0.016	*** 0.016	*** 0.016	*** 0.016	*** 0.013	*** 0.013	*** 0.020	*** 0.020
	(0.002)	(0.004)	(0.005)	(0.005)	(0.003)	(0.004)	(0.004)	(0.004)	(0.002)	(0.004)	(0.005)	(0.005)
σ	*** 0.050	*** 0.050	*** 0.050	*** 0.050	*** 0.041	*** 0.041	*** 0.041	*** 0.041	*** 0.050	***0.050	*** 0.052	*** 0.052
	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)
ϕ			0.002	0.002			0.000	0.000				
			(0.003)	(0.003)			(0.001)	(0.001)				
κ_2			* 0.104	* 0.104			0.004	0.004				
			(0.046)	(0.049)			(0.007)	(0.007)				
υ					*** 0.051	*** 0.051	*** 0.051	*** 0.051				
					(0.006)	(0.006)	(0.006)	(0.006)				
κ_3					· 0.017	· 0.017	· 0.017	· 0.017				
					(0.010)	(0.010)	(0.010)	(0.010)				
ζ									0.035	0.033	* 0.012	* 0.012
									(2435921)	(531611)	(0.006)	(0.006)
ℓ_0									0.000	0.000	*** 1.757	*** 1.757
									(0.096)	(0.096)	(0.265)	(0.267)
ℓ_1											*** 0.564	*** 0.564
											(0.090)	(0.095)
κ_4											* 0.084	* 0.084
											(0.039)	(0.039)
ϵ		0.000		0.000		0.000		0.000		0.000		0.000
		(0.003)		(0.003)		(0.003)		(0.004)		(0.003)		(0.004)
L	275.19	275.19	277.03	277.03	285.59	285.59	285.67	285.67	275.19	275.19	349.94	349.94

Table 5: Parameter estimates of the nested models for Australia (“AU”). We compute these estimates by maximizing the simulated likelihood $\hat{L}_m^K(\theta)$ of Algorithm 3.1 with $K = 1000$ subject to the restrictions in Table 2. We use the Nelder-Mead algorithm in R to compute the maximizer of the simulated likelihood. Given in parenthesis are asymptotic standard errors computed as in Guay and Schwenkler (2018). 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.

	DE											
	RW	RW+ME	LRGR	LRGR+ME	LRVR	LRVR+ME	LRR	LRR+ME	DIS	DIS+ME	TVDIS	TVDIS+ME
μ	*** 0.021	*** 0.021	*** 0.021	*** 0.021	*** 0.021	*** 0.021	*** 0.017	*** 0.017	*** 0.021	*** 0.021	*** 0.024	*** 0.024
	(0.003)	(0.003)	(0.005)	(0.005)	(0.001)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)
σ	*** 0.052	*** 0.052	*** 0.052	*** 0.052	*** 0.038	*** 0.038	*** 0.037	*** 0.037	*** 0.052	*** 0.052	*** 0.055	*** 0.055
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)
ϕ			0.000	0.000			0.000	0.000				
			(0.001)	(0.001)			(0.000)	(0.000)				
κ_2			** 0.036	** 0.036			*** 0.023	*** 0.023				
			(0.011)	(0.011)			(0.006)	(0.006)				
υ					*** 0.119	*** 0.120	*** 0.126	*** 0.128				
					(0.008)	(0.008)	(0.009)	(0.009)				
κ_3					** 0.053	** 0.053	** 0.054	** 0.054				
					(0.015)	(0.016)	(0.016)	(0.015)				
ζ									0.010	0.031	** 0.038	** 0.038
									(554993)	(7818629)	(0.011)	(0.011)
ℓ_0									0.000	0.000	*** 0.758	*** 0.758
									(0.080)	(0.080)	(0.096)	(0.097)
ℓ_1											*** 0.448	*** 0.448
											(0.038)	(0.038)
κ_4											** 0.092	** 0.092
											(0.035)	(0.035)
ϵ		0.000		0.000		0.001		0.001		0.000		0.000
		(0.002)		(0.003)		(0.001)		(0.002)		(0.002)		(0.003)
L	395.99	395.99	397.04	397.04	417.56	417.79	422.98	423.29	395.99	395.99	484.35	484.35

Table 6: Parameter estimates of the nested models for Germany ("DE"). We compute these estimates by maximizing the simulated likelihood $\hat{L}_m^K(\theta)$ of Algorithm 3.1 with $K = 1000$ subject to the restrictions in Table 2. We use the Nelder-Mead algorithm in R to compute the maximizer of the simulated likelihood. Given in parenthesis are asymptotic standard errors computed as in Guay and Schwenkler (2018). 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.

	JP											
	RW	RW+ME	LRGR	LRGR+ME	LRVR	LRVR+ME	LRR	LRR+ME	DIS	DIS+ME	TVDIS	TVDIS+ME
μ	*** 0.023	*** 0.025	*** 0.039	*** 0.039	*** 0.023	*** 0.023	*** 0.019	*** 0.019	*** 0.025	*** 0.025	*** 0.032	*** 0.032
	(0.006)	(0.006)	(0.003)	(0.004)	(0.001)	(0.003)	(0.003)	(0.002)	(0.006)	(0.006)	(0.003)	(0.003)
σ	*** 0.075	*** 0.074	*** 0.075	*** 0.074	*** 0.038	*** 0.038	*** 0.039	*** 0.039	*** 0.075	*** 0.074	*** 0.076	*** 0.076
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
ϕ			*** 0.005	*** 0.005			0.001	* 0.001				
			(0.001)	(0.001)			(0.000)	(0.000)				
κ_2			· 0.031	· 0.031			0.015	* 0.015				
			(0.017)	(0.015)			(0.011)	(0.008)				
υ					*** 0.143	*** 0.145	*** 0.138	*** 0.140				
					(0.006)	(0.011)	(0.006)	(0.006)				
κ_3					* 0.056	* 0.056	** 0.058	* 0.058				
					(0.016)	(0.022)	(0.024)	(0.026)				
ζ									0.000	0.000	0.001	0.017
									(19278322)	(5663794)	(0.021)	(0.018)
ℓ_0									0.000	0.000	*** 0.891	*** 0.892
									(0.086)	(0.086)	(0.179)	(0.180)
ℓ_1											*** 0.421	*** 0.421
											(0.061)	(0.061)
κ_4											** 0.084	** 0.085
											(0.029)	(0.030)
ϵ		** 0.006		* 0.006		0.002		*** 0.002		** 0.006		0.000
		(0.002)		(0.003)		(0.002)		(0.001)		(0.002)		(0.003)
L	299.53	299.94	301.62	302.99	334.53	334.93	336.44	336.91	299.49	299.94	369.76	369.76

Table 7: Parameter estimates of the nested models for Japan ("JP"). We compute these estimates by maximizing the simulated likelihood $\hat{L}_m^K(\theta)$ of Algorithm 3.1 with $K = 1000$ subject to the restrictions in Table 2. We use the Nelder-Mead algorithm in R to compute the maximizer of the simulated likelihood. Given in parenthesis are asymptotic standard errors computed as in Guay and Schwenkler (2018). 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	RW+ME	LRGR	LRGR+ME	LRVR	LRVR+ME	LRR	LRR+ME	DIS	DIS+ME	TVDIS	TVDIS+ME
μ	*** 0.015	*** 0.015	*** 0.015	*** 0.015	*** 0.015	*** 0.015	*** 0.015	*** 0.015	*** 0.015	*** 0.015	*** 0.020	*** 0.020
	(0.001)	(0.003)	(0.002)	(0.003)	(0.001)	(0.001)	(0.003)	(0.003)	(0.001)	(0.003)	(0.002)	(0.002)
σ	*** 0.040	*** 0.040	*** 0.040	*** 0.040	*** 0.041	*** 0.041	*** 0.041	*** 0.041	*** 0.040	*** 0.040	*** 0.038	*** 0.038
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
ϕ			0.000	0.000			0.000	0.000				
			(0.001)	(0.001)			(0.001)	(0.001)				
κ_2			· 0.039	* 0.047			** 0.048	** 0.047				
			(0.021)	(0.024)			(0.017)	(0.017)				
υ					0.005	0.005	0.008	0.008				
					(0.008)	(0.008)	(0.010)	(0.010)				
κ_3					* 0.043	* 0.043	** 0.043	** 0.043				
					(0.013)	(0.013)	(0.020)	(0.023)				
ζ									0.032	0.029	0.038	0.038
									(801046)	(1475564)	(0.028)	(0.029)
ℓ_0									0.000	0.000	*** 0.758	*** 0.758
									(0.074)	(0.074)	(0.115)	(0.121)
ℓ_1											*** 0.466	*** 0.466
											(0.051)	(0.052)
κ_4											*** 0.106	*** 0.106
											(0.021)	(0.021)
ϵ		0.001		0.001		0.001		0.001		0.001		0.000
		(0.002)		(0.002)		(0.001)		(0.002)		(0.002)		(0.001)
L	518.58	518.70	521.05	521.27	520.02	520.19	522.79	522.94	518.58	518.70	625.02	625.04

Table 8: Parameter estimates of the nested models for the United States ("US"). We compute these estimates by maximizing the simulated likelihood $\hat{L}_m^K(\theta)$ of Algorithm 3.1 with $K = 1000$ subject to the restrictions in Table 2. We use the Nelder-Mead algorithm in R to compute the maximizer of the simulated likelihood. Given in parenthesis are asymptotic standard errors computed as in Guay and Schwenkler (2018). 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.

	AU				
	Data	RW+ME	LRR+ME	TVDIS+ME	FULL+ME
Mean	0.014	0.016	0.019	0.011	0.010
Standard deviation	0.050	0.053	0.045	0.044	0.075
Skewness	−0.982	0.070	0.060	−0.030	−0.278
Kurtosis	7.910	2.669	2.869	3.479	1.726
Probability of a decline of more than 5%	0.074	0.109	0.061	0.098	0.268
Probability of a decline of more than 10%	0.028	0.009	0.003	0.008	0.079
First-order autocovariance	0.000	0.000	−0.001	0.001	0.001

Table 9: Model-implied moments of Australian consumption growth for the nested models. We compute model-implied moments using the methodology of Guay and Schwenkler (2018) with $K = 10^5$. We compare these model-implied moments to the corresponding moments in the data. All moments are measured in annual terms.

	DE				
	Data	RW+ME	LRR+ME	TVDIS+ME	FULL+ME
Mean	0.017	0.021	0.017	−0.004	−0.004
Standard deviation	0.053	0.057	0.041	0.070	0.071
Skewness	−0.550	−0.023	−0.006	−0.167	−0.126
Kurtosis	7.264	2.936	3.035	2.936	2.865
Probability of a decline of more than 5%	0.057	0.109	0.050	0.241	0.249
Probability of a decline of more than 10%	0.038	0.001	0.000	0.026	0.027
First-order autocovariance	0.001	0.001	−0.001	0.005	0.005

Table 10: Model-implied moments of German consumption growth for the nested models. We compute model-implied moments using the methodology of Guay and Schwenkler (2018) with $K = 10^5$. We compare these model-implied moments to the corresponding moments in the data. All moments are measured in annual terms.

	JP				
	Data	RW+ME	LRR+ME	TVDIS+ME	FULL+ME
Mean	0.022	0.025	0.021	0.020	0.022
Standard deviation	0.069	0.069	0.043	0.063	0.062
Skewness	−1.509	−0.020	0.043	−0.029	−0.015
Kurtosis	19.887	2.435	2.901	2.576	2.589
Probability of a decline of more than 5%	0.059	0.149	0.048	0.144	0.137
Probability of a decline of more than 10%	0.022	0.031	0.002	0.025	0.023
First-order autocovariance	0.001	0.001	−0.001	0.001	0.001

Table 11: Model-implied moments of Japanese consumption growth for the nested models. We compute model-implied moments using the methodology of Guay and Schwenkler (2018) with $K = 10^5$. We compare these model-implied moments to the corresponding moments in the data. All moments are measured in annual terms.

	US				
	Data	RW+ME	LRR+ME	TVDIS+ME	FULL+ME
Mean	0.015	0.015	0.015	0.002	0.018
Standard deviation	0.038	0.043	0.044	0.047	0.045
Skewness	−0.071	0.121	0.131	−0.191	−0.117
Kurtosis	3.528	3.000	2.879	2.932	3.154
Probability of a decline of more than 5%	0.038	0.062	0.065	0.147	0.063
Probability of a decline of more than 10%	0.005	0.001	0.001	0.019	0.009
First-order autocovariance	0.000	0.002	0.002	0.001	0.001

Table 12: Model-implied moments of American consumption growth for the nested models. We compute model-implied moments using the methodology of Guay and Schwenkler (2018) with $K = 10^5$. We compare these model-implied moments to the corresponding moments in the data. All moments are measured in annual terms.

Unconstrained model	Constrained model		AU	DE	JP	US
LRGR+ME	RW+ME	Statistic	3.68	2.10	* 6.10	· 5.14
		Degrees of freedom	2	2	2	2
		p-value	0.160	0.350	0.047	0.077
LRVR+ME	RW+ME	Statistic	*** 20.80	*** 43.60	*** 69.98	2.98
		Degrees of freedom	2	2	2	2
		p-value	0.000	0.000	0.000	0.225
LRR+ME	RW+ME	Statistic	*** 20.96	*** 54.60	*** 73.94	· 8.47
		Degrees of freedom	4	4	4	4
		p-value	0.000	0.000	0.000	0.076
DIS+ME	RW+ME	Statistic	0.000	0.000	0.000	0.000
		Degrees of freedom	2	2	2	2
		p-value	1.000	1.000	1.000	1.000
TVDIS+ME	RW+ME	Statistic	*** 149.50	*** 176.74	*** 139.64	*** 212.667
		Degrees of freedom	4	4	4	4
		p-value	0.000	0.000	0.000	0.000
FULL+ME	TVDIS+ME	Statistic	5.56	5.28	6.34	*** 13.45
		Degrees of freedom	4	4	4	3
		p-value	0.235	0.260	0.175	0.009
FULL+ME	LRR+ME	Statistic	*** 134.10	*** 127.42	*** 72.04	*** 217.65
		Degrees of freedom	4	4	4	4
		p-value	00.00	0.000	0.000	0.000

Table 13: Likelihood ratio tests of the null hypothesis that an unconstrained model fits the data as well as a constrained model. The test statistic is given in (16). 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. *** 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.

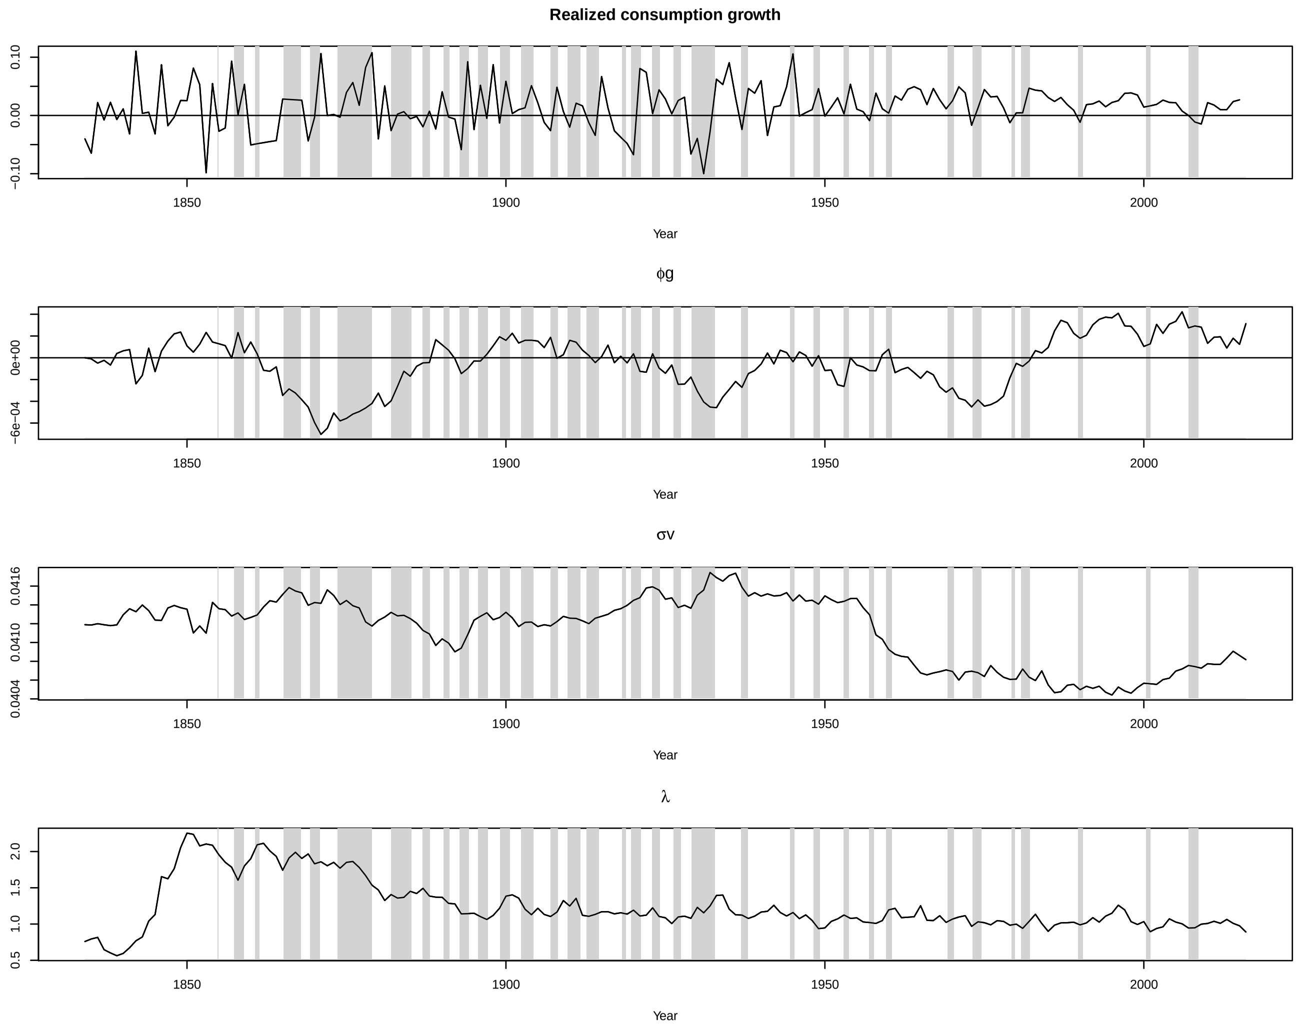

Figure 1: Filtered sample path of the latent long-run growth risk factor ($\mathbb{E}[\phi g_t \,|\, \mathbf{D}_t]$), stochastic volatility ($\mathbb{E}[\sigma v_t \,|\, \mathbf{D}_t]$), and disaster intensity ($\mathbb{E}[\lambda_t \,|\, \mathbf{D}_t]$) for the United States. We evaluate these filters using the methodology of Guay and Schwenkler (2018) with $K = 10^7$ and the parameter estimates for the model "FULL+ME" from Table 3. We also show the NBER recession periods in grey

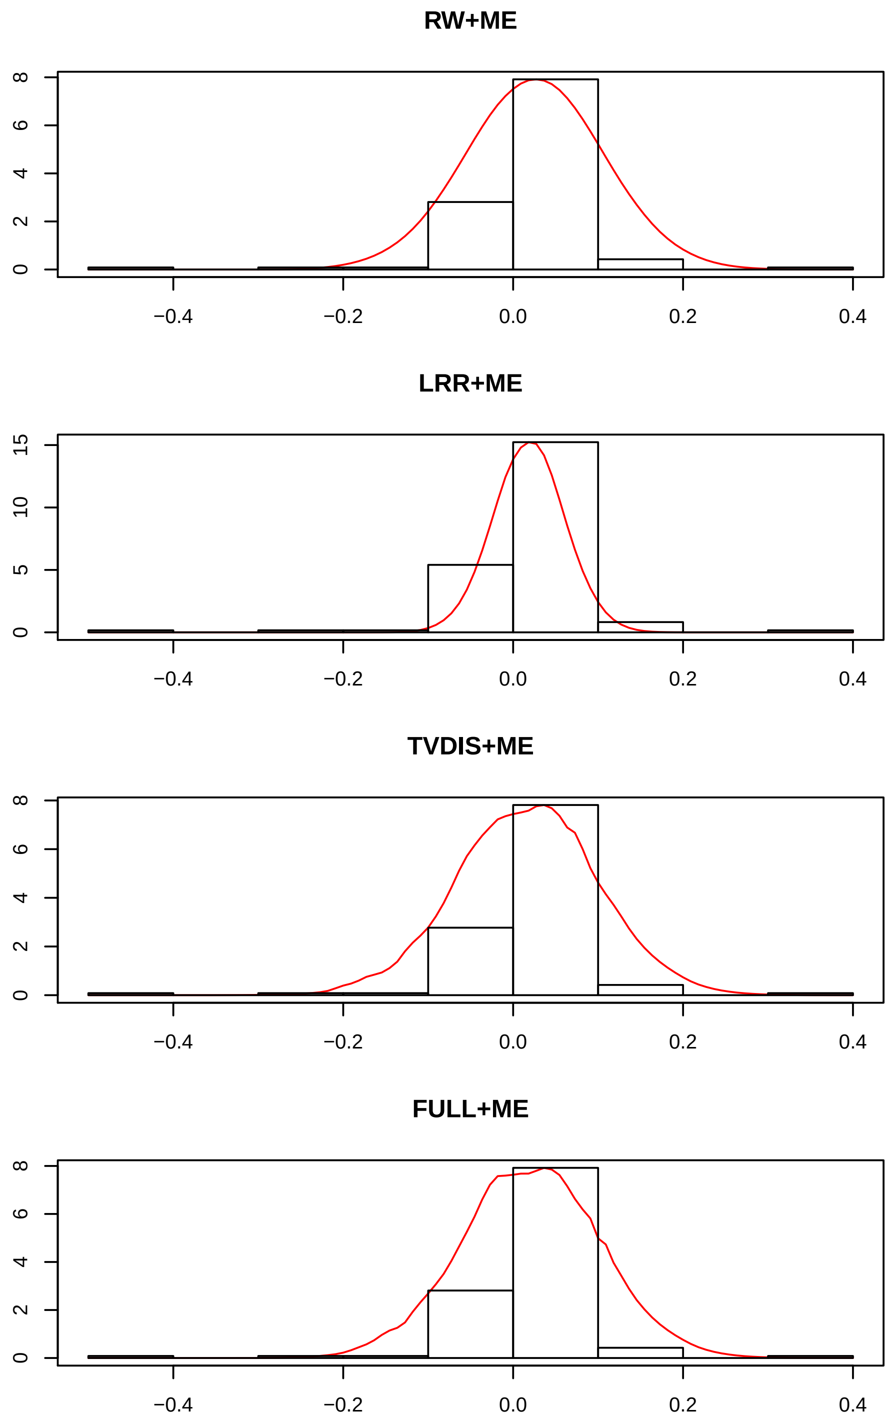

Figure 2: *Model-implied marginal distributions for Japan.* This figure shows in red the implied marginal distributions of consumption growth in Japan for the different models we consider. It also shows histograms of realized consumption growth in black. The marginal distributions are computed using the methodology of Guay and Schwenkler (2018) with $K = 10^5$ with the parameter estimates from Tables 3 and 7.

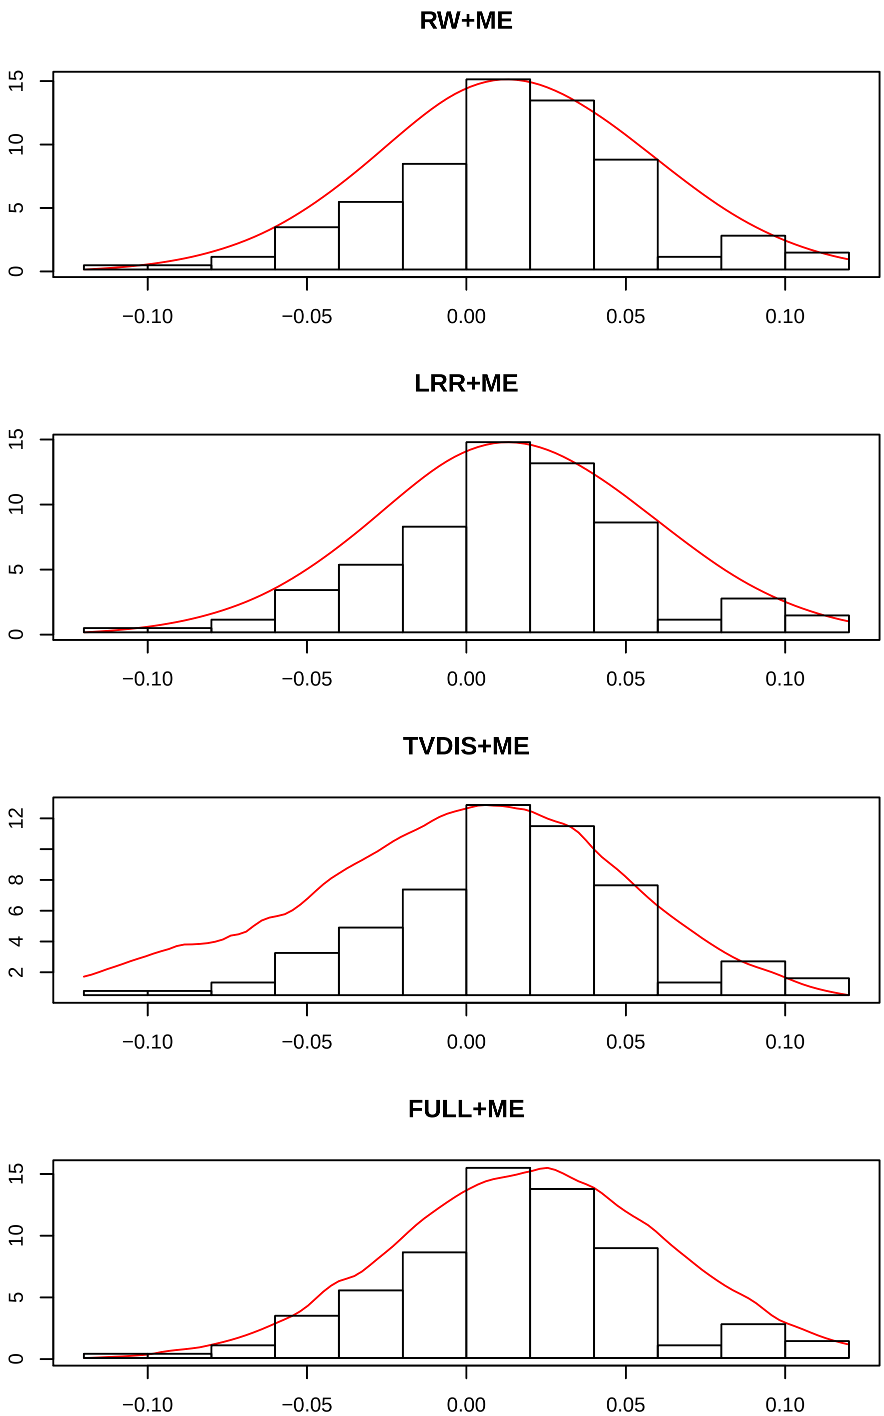

Figure 3: *Model-implied marginal distributions for the United States.*. This figure shows in red the implied marginal distributions of consumption growth in the U.S. for the different models we consider. It also shows histograms of realized consumption growth in black. The marginal distributions are computed using the methodology of Guay and Schwenkler (2018) with $K = 10^5$ with the parameter estimates from Tables 3 and 8.