Rare events and long-run risks ☆

Robert J. Barro [a,1], Tao Jin [b,*]

[a] *Harvard University, USA*
[b] *Tsinghua University, China*

ARTICLE INFO

ABSTRACT

Rare events (RE) and long-run risks (LRR) are complementary approaches for characterizing macroeconomic variables and understanding asset pricing. We estimate a model with RE and LRR using long-term consumption data for 42 economies, identify these two types of risks simultaneously from the data, and reveal their distinctions. RE typically associates with major historical episodes, such as world wars and depressions and analogous country-specific events. LRR reflects gradual processes that influence long-run growth rates and volatility. A match between the model and observed average rates of return on equity and short-term bonds requires a coefficient of relative risk aversion, γ, around 6. Most of the explanation for the equity premium derives from RE although LRR makes a moderate contribution. However, LRR helps in fitting the Sharpe ratio. Generating good matches to the equity premium and Sharpe ratio simultaneously is still challenging.

© 2020 Elsevier Inc. All rights reserved.

Rare macroeconomic events, denoted RE, provide one approach for modeling the long-term evolution of macroeconomic variables such as GDP and consumption. Another approach, called long-run risks or LRR, emphasizes variations in the long-run growth rate and the variance of shocks to this growth rate (stochastic volatility). An extensive literature has studied RE and LRR as distinct phenomena, but a joint approach does better at describing the macro data. Moreover, although we prefer a model that incorporates both features, we can assess the relative contributions of RE and LRR for explaining asset-pricing properties, such as the average equity premium and the volatility of equity returns.

As in previous research, this study treats RE and LRR as latent variables. Our formalization of the distinct features of RE and LRR allows us to isolate these two forces using data on real per capita consumer expenditure for 42 economies going

back as far as 1851 and ending in 2012 (4 14 country-year observations). The estimated model indicates that R_E comprises *sporadic, drastic, and jumping outbursts*, whereas RR exhibits *persistent, moderate, and smooth fluctuations*.

With respect to R_E our results include characterizations for when the world and individual countries are in disaster states and by how much. We also isolate patterns of economic recovery, related to the extent to which disaster shocks have permanent or temporary impacts. At the world level, the periods labeled as R_E (based on posterior probability distributions) correspond to familiar historical events, such as the world wars, the Great Depression, and possibly the Great Influenza epidemic of 1918–20 (but not the recent Great Recession). For individual or small groups of countries, examples of events associated with rare disasters are the Asian Financial Crisis of 1997-8, the Russian Revolution and Civil War after World War I, the 1973 Chilean coup and its aftermath, and the German hyperinflation in 1921-24.

Similarly, for RR, our results include ex-post characterizations of movements in the long-run growth rate and volatility. In contrast to R_E RR exhibits much smoother, low-frequency evolution. For example, for the United States, the long-run growth component is estimated to be well above normal for 1962-67, 1971, 1983-85, and 1997-98 —recent periods typically viewed as favorable for economic growth. At earlier times, the long-run growth rate is unusually high in 1933-36 (recovery from the Great Depression), 1919, and 1875-79 (resumption of the gold standard). On the down side, the estimated U.S. long-run growth rate is unusually low in 2007-09 (Great Recession), 1990, 1979, 1910-13, 1907, 1892-93, 1859-65, and 1853-55.

As examples for other countries, the estimated long-run growth rate is high in Germany for 1945-71; Japan for 1945-72; Chile for 1986-96, 2003-06, and 2009-11; Russia for 1999-2011; and the United Kingdom for 1983-88 and 1995-2002. Weak periods for the long-run growth rate include Russia in 1989-97 and the United Kingdom in 2007-11.

The estimated process for stochastic volatility is even smoother than that for the long-run growth rate. The results for recent years exhibit the frequently mentioned pattern of moderation—the estimated volatility was particularly low in the late 1990s for many countries, including the United States, Germany, and Japan. In contrast, Russia experienced a sharp rise in volatility from 1973 to 2007.

To assess asset pricing, we embed the estimated time-series process for consumption into an endowment economy with a representative agent that has Epstein-Zin-Weil (EZW) preferences (Epstein and Zin (1989) and Weil (1990)). This analysis generates predictions for the average equity premium, the volatility of equity returns, and so on. Then we compare these predictions with averages found in the long-term data for a group of countries.

The rest of the paper is organized as follows. Section 1 relates our study to the previous literature on rare macroeconomic events and long-run risks. Section 2 lays out our formal model, which includes rare events (partly temporary, partly permanent) and long-run risks (including stochastic volatility). Section 3 discusses the long-term panel data on consumer expenditure, describes our method of estimation, and presents empirical results related to R_E RR, the distinctions between them, and the time evolution of consumer spending in each country. The analysis includes a detailed description for six illustrative countries of the evolution of posterior means of the key variables related to rare events and long-run risks. Section 4 presents the framework for asset pricing. We draw out the implications of the estimated processes for consumer spending for various statistics, including the average equity premium, the volatility of equity returns, and the Sharpe Ratio. Section 5 discusses the potential addition of time variation in the rare-disaster probability. Section 6 has conclusions.

1. Relation to the literature

The notion of rare macroeconomic events has been employed by researchers to explain a variety of phenomena in asset and foreign-exchange markets, as surveyed in Barro and Ursúa (2012). Examples of this literature are Gabaix (2012), Gourio (2008, 2012), Farhi and Gabaix (2016), Farhi et al. (2015), Wachter (2013), Seo and Wachter (2016), and Colacito and Croce (2013).

Bansal and Yaron (2004), henceforth BY, introduced the idea of long-run risks. The central notion is that small but persistent shocks to expected growth rates and to the volatility of shocks to growth rates are important for explaining various asset-market phenomena, including the high average equity premium and the high volatility of stock returns. The main results in BY and in the updated study by Bansal et al. (2010) required a coefficient of relative risk aversion, γ, around 10, much higher than the values needed in the rare-disasters literature. (BY assumed an intertemporal elasticity of substitution of 1.5 and also assumed substantial leverage in the relation between dividends and consumption.) In our study, we incorporate the long-run risks framework of BY, along with an updated specification for rare macroeconomic events.

The idea of long-run risks has been applied to many aspects of asset and foreign-exchange markets. This literature includes Bansal and Shaliastovich (2013); Bansal et al. (2005); Hansen et al. (2008); Malloy et al. (2009); Croce et al. (2015); Chen (2010); Colacito and Croce (2011); and Nakamura et al. (2017). Beeler and Campbell (2012) provide a critical empirical evaluation of the long-run-risks model.

There is a large literature investigating separately the implications of rare events, RE, and long-run risks, LRR. However, our view is that—despite the order-of-magnitude increase in the required numerical analysis—it is important to assess the two core ideas, RE and LRR, in a simultaneous manner.[2] This study reports the findings from this joint analysis.

2. Model of rare events and long-run risks

The model allows for rare events, RE, and long-run risks, LRR. The RE part follows Nakamura et al. (2013) (or NSBU) in allowing for macroeconomic disasters of stochastic size and duration, along with recoveries that are gradual and of stochastic proportion. We modify the NSBU framework in various dimensions, including the specification of probabilities for world and individual country transitions between normal and disaster states. Most importantly, we expand on NSBU by incorporating long-run risks, along the lines of Bansal and Yaron (2004). The LRR specification allows for fluctuations in long-run growth rates and for stochastic volatility.

2.1. Components of consumption

As in NSBU, the log of consumption per capita for country i at time t, c_{it}, is the sum of three unobserved variables:

$$c_{it} = x_{it} + z_{it} + \sigma_{\varepsilon i}\varepsilon_{it},\tag{1}$$

where x_{it} is the "potential level" (or permanent part) of the log of per capita consumption and z_{it} is the "event gap," which describes the deviation of c_{it} from its potential level due to current and past rare events. The potential level of consumption and the event gap depend on the disaster process, as detailed below. The term $\sigma_{\varepsilon i}\varepsilon_{it}$ is the error term, where ε_{it} is an i.i.d. standard normal variable. The standard deviation, $\sigma_{\varepsilon i}$, of the error term varies by country. We also allow $\sigma_{\varepsilon i}$ to take on two values for each country, one up to 1945 and another thereafter.[3] This treatment allows for post-WWII moderation in observed consumption volatility particularly because of improved measurement in national accounts—see Romer (1986) and Balke and Gordon (1989). In this study, we view $\sigma_{\varepsilon i}\varepsilon_{it}$ as measurement error, rather than a consumption shock. Thus, it is attributed to neither rare disasters nor long-run risks.

2.2. Disaster probabilities

We follow NSBU, but with significant modifications, in assuming that rare macroeconomic events involve disaster and normal states. Each state tends to persist over time, but there are possibilities for transitioning from one state to the other. The various probabilities have world and country-specific components.

For the world component, we have in mind the influence from major international catastrophes such as the two world wars and the Great Depression of the early 1930s. Additional possible examples are the Great Influenza Epidemic of 1918-20, the threat from climate change, and the current Coronavirus Pandemic.[4] However, the recent global financial crisis of 2008-09 turns out not to be sufficiently important to show up as a world disaster.

We characterize the world process with two probabilities—one, denoted p_0, is the probability of moving from normalcy to a global disaster state (such as the start of a world war or global depression), and two, denoted p_1, is the probability of

---

[2] Nakamura et al. or NSS (2017, section 3) filter the consumption data for crudely estimated disaster effects based on the results in Nakamura et al. (2013) or NSBU. Thus, NSS do not carry out a joint analysis of rare events and long-run risks. This joint analysis was also not in NSBU, which neglected long-run risks. In their analysis of asset pricing, NSS consider only the role of long-run risks (applied to their disaster-filtered data), whereas NSBU allowed only for effects from rare events. Thus, neither NSS nor NSBU carried out a joint analysis of rare events and long-run risks.

[3] When the data for country i begin after 1936, $\sigma_{\varepsilon i}$ takes on only one value.

[4] See Barro (2015) for an application of the rare-events framework to environmental issues. See Barro et al. (2020) for an analysis of the ongoing coronavirus pandemic as a realization of a rare disaster.

2.6. Stochastic volatility

Stochastic volatility, σ_{it}, enters in equations (4) and (5). We follow Bansal and Yaron (2004, p. 1487) in modeling the evolution of volatility as an AR(1) process for the variance:

$$\sigma_{it}^2 = \sigma_i^2 + \rho_\sigma \left(\sigma_{i,t-1}^2 - \sigma_i^2 \right) + \sigma_{\omega i} \omega_{it}, \tag{6}$$

where σ_i^2 is the average country-specific variance, and ρ_σ is a first-order autoregressive coefficient, with $0 \le \rho_\sigma < 1$. The shock includes the standard normal variable ω_{it} multiplied by the country-specific volatility of volatility, $\sigma_{\omega i}$. In the estimation, we use a method similar to Bansal and Yaron (2004, p. 1495, n. 13) in constraining σ_{it}^2 to be non-negative (see Appendix A.3).

2.7. Dynamics of event gaps

Returning to equation (1), we now consider the event gap, z_{it}, which describes the deviation of c_{it} from its potential level due to current and past rare events. We assume, following NSBU, that z_{it} follows a modified autoregressive process:

$$z_{it} = \rho_z z_{i,t-1} + I_{it}\phi_{it} - I_{it}\eta_{it} + \sigma_{vi}v_{it}, \tag{7}$$

where ρ_z is a first-order autoregressive coefficient, with $0 \le \rho_z < 1$. The term $I_{it}\phi_{it}$ picks up the immediate effect of a disaster on consumption, whereas the term $I_{it}\eta_{it}$ captures the permanent part of this effect. Thus, the term $I_{it} \cdot (\phi_{it} - \eta_{it})$ is the temporary part of the disaster shock. The error term includes the standard normal variable v_{it} multiplied by the country-specific constant volatility σ_{vi}.

The direct effect of a disaster during period t appears in equation (7) as the term $I_{it}\phi_{it}$. We assume that ϕ_{it} is negative, and we model it as a truncated normal distribution (with mean and variance for the non-truncated distribution that are constant over time and across countries). Thus, in the short run, a disaster lowers c_{it} in equation (1). However, as the event gap vanishes in equation (7), part of this disaster effect on c_{it} disappears. Specifically, for given $I_{it}\eta_{it}$, the shock $I_{it}\phi_{it}$ does not affect c_{it} in the long run.

The long-run impact of a disaster involves the term $-I_{it}\eta_{it}$ in equation (7), which operates in conjunction with the term $+I_{it}\eta_{it}$ in equation (4). The combination of these two terms means that the short-run effect of η_{it} on c_{it} in equation (1) is nil. However, as the event gap, z_{it}, vanishes, the long-run impact on consumption approaches η_{it}. Thus, if $\eta_{it} < 0$ (the typical case), the effect on the long-run consumption level is negative.

If $\eta_{it} = \phi_{it}$, the long- and short-run effects of a disaster coincide; that is, disasters have only permanent effects on c_{it}. If $\eta_{it} = 0$, the long-run effect of a disaster is nil; that is, disasters have only temporary effects on c_{it}. We find empirically, as do NSBU, that recoveries tend to occur but are typically only partial. This result corresponds to a mean for η_{it} that is negative but smaller in magnitude than that for ϕ_{it}.

2.8. Consumption growth

The estimation is based on the observable growth rate of per capita consumption, Δc_{it} (based on the available data on personal consumer expenditure). To see how this variable relates to the underlying rare events and long-run risks, start by taking a first-difference of equation (1). Then substitute for Δx_{it} from equation (4) and for z_{it} and $z_{i,t-1}$ from equation (7) to get:

$$\Delta c_{it} = \underbrace{I_{it}\phi_{it} - (1-\rho_z)I_{i,t-1}\phi_{i,t-1} + (1-\rho_z)I_{i,t-1}\eta_{i,t-1} - \rho_z(1-\rho_z)z_{i,t-2}}_{RE}$$

$$+ \underbrace{\mu_i + \chi_{i,t-1}}_{\text{long-run growth rate}} + \text{error term}. \tag{8}$$

Equation (8) shows that consumption growth can be decomposed into a rare-events (RE) component, the long-run growth rate (which includes the persistent component of the consumption growth, the main part of the LRR), and the error term. This error depends on u_{it} (equation (4)) and the contemporaneous and lagged values of ε_{it} (equation (1)) and v_{it} (equation (7)).

To bring out the main properties for the RE term, assume first that $\rho_z = 0$ in equation (8), so that event gaps have zero persistence over time in equation (7). In an RE state ($I_{it} = 1$), the shock $\phi_{it} < 0$ gives the initial downward effect on consumption growth. For given η_{it}, this effect exactly reverses the next period—that is, the effect on the level of c is temporary, so that an equal-size rise in consumption growth follows the initial fall. In contrast, if $\eta_{it} = \phi_{it}$, the effect on the level of c is permanent, and there is no impact on next period's consumption growth rate. The lagged term $z_{i,t-2}$ in equation (8) brings in more lags of rare-events shocks through the dynamics of event gaps in equation (7). This lag structure applies when $\rho_z \ne 0$.

To assess RR, consider the term for the long-run growth rate in equation (6). The first part, μ_i, is assumed to be constant for country i. The RR effect is mainly given by $\chi_{i,t-1}$, which is the variable part of the long-run growth rate. This term evolves in accordance with equations (5) and (6), which allow for stochastic volatility.

2.9. Alternative decomposition of consumption growth

The previous decomposition focuses on the roles of RE and the long-run growth rate, the main part of RR. The shock that includes stochastic volatility, $\sigma_{i,t-1}u_{it}$, does not show up there explicitly. However, we can decompose the consumption growth rate in a different way to separate the term $\sigma_{i,t-1}\eta_{it}$ from the other error terms.

For country i, define the *consumption growth gap* $\widetilde{\Delta c}_{it}$ as the difference between the actual and long-term average growth rate μ_i:

$$\widetilde{\Delta c}_{it} \triangleq \Delta c_{it} - \mu_i = c_{it} - c_{i,t-1} - \mu_i.$$

This growth rate can be decomposed into four components as follows:

$$\widetilde{\Delta c}_{it} \triangleq RE_{it} + \chi_{i,t-1} + \sigma_{i,t-1}u_{it} + N_{it},$$

where

$$RE_{it} = I_{it}\eta_{it} + \Delta z_{it} = I_{it}\eta_{it} + z_{it} - z_{i,t-1}$$

and

$$N_{it} = \Delta(\sigma_{\varepsilon i}\varepsilon_{it}) = \sigma_{\varepsilon i}\varepsilon_{it} - \sigma_{\varepsilon i}\varepsilon_{i,t-1}.$$

The RE_{it} term is basically the same as the RE component defined in Section 2, except that RE_{it} contains the shocks ν_{it} and $\nu_{i,t-1}$. The slow-varying component $\chi_{i,t-1}$ characterizes the long-run growth rate, and N_{it} is the noise or measurement-error term. The long-term mean values of $\chi_{i,t-1}$, $\sigma_{i,t-1}u_{it}$, and N_{it} are 0, while that of RE_{it} is not. Let RE_{it}^{DM} denote the *demeaned RE_{it}*, and

$$\Delta c_{it}^{DM} \triangleq RE_{it}^{DM} + \chi_{i,t-1} + \sigma_{i,t-1}u_{it} + N_{it}$$

denote the *demeaned consumption growth gap*. The terms in this last decomposition will be identified after the model is estimated (see Section 3.2).

3. Data, estimation method, and empirical results

We use an expanded version of the data on annual consumption (real per capita personal consumer expenditure) provided for 42 economies in Barro and Ursúa (2010). We extended on these data by including observations as far back as 1851 (rather than 1870) and going through 2012. There are 4714 country-year observations. Appendix A.1 provides details.

We follow NSBU in estimating the model with the Bayesian Markov-Chain Monte-Carlo (MCMC) method. RE and RR are shocks of different nature, and the statistical distinctions between them enable us to identify them. Bayesian MCMC is an appropriate choice for estimating the model because, first, it is a standard and widely adopted estimation method; second, the necessary identifying information can be conveniently incorporated into prior beliefs; and, third, it is relatively easy to implement for as complicated a model as the one proposed here.[6] Our implementation of Bayesian MCMC features nearly flat prior distributions for the various underlying parameters. See Appendix A.3 for details. Here, we focus on the posterior means of each parameter.

3.1. Estimated model

Table 1 contains the posterior means and standard deviations for the main parameters of the model. These parameters apply across countries and over time.

1. Transition probabilities. The first group of parameters in Table 1 applies to transition probabilities between normal and disaster states. With respect to a world event, we find that p_0, the estimated probability of moving from a normal to a disaster state, is 2.9% per year. Once entering a disaster, there is a lot of persistence: the estimated conditional probability, p_1, of the world remaining in a disaster state the following year is 65.8 %.

The probability of a disaster for an individual country depends heavily on the global situation and also on whether the country was in a disaster state in the previous year. If there is no contemporaneous world disaster, the estimated probability,

---

[6] Bansal et al. (2016) propose a method to estimate the RR model with time aggregation using the Generalized Method of Moments (GMM). However, that method is not helpful in our setting because we are using annual data, and the decision interval of the agents in Bansal et al. (2016) is only 33 days. See also notes 12 and 13.

Table 1
Estimated parameters—model with rare events and long-run risks.

Parameter	Definition	Posterior mean	Posterior s.d.	5% & 95% Percentiles
	World disaster probability, conditional on:			
p_0	No prior-year world disaster	0.029	0.011	0.012, 0.047
p_1	Prior-year world disaster	0.658	0.139	0.397, 0.854
	Country disaster probability, conditional on:			
q_{00}	No prior-year disaster, no current world disaster	0.0066	0.0022	0.0035, 0.0107
q_{10}	Prior-year disaster, no current world disaster	0.719	0.050	0.632, 0.790
q_{01}	No prior-year disaster, current world disaster	0.360	0.052	0.304, 0.470
q_{11}	Prior-year disaster, current world disaster	0.857	0.037	0.782, 0.897
	Parameters that are constant across countries:			
ρ_z	AR(1) coefficient for event gap (eq. (7)			
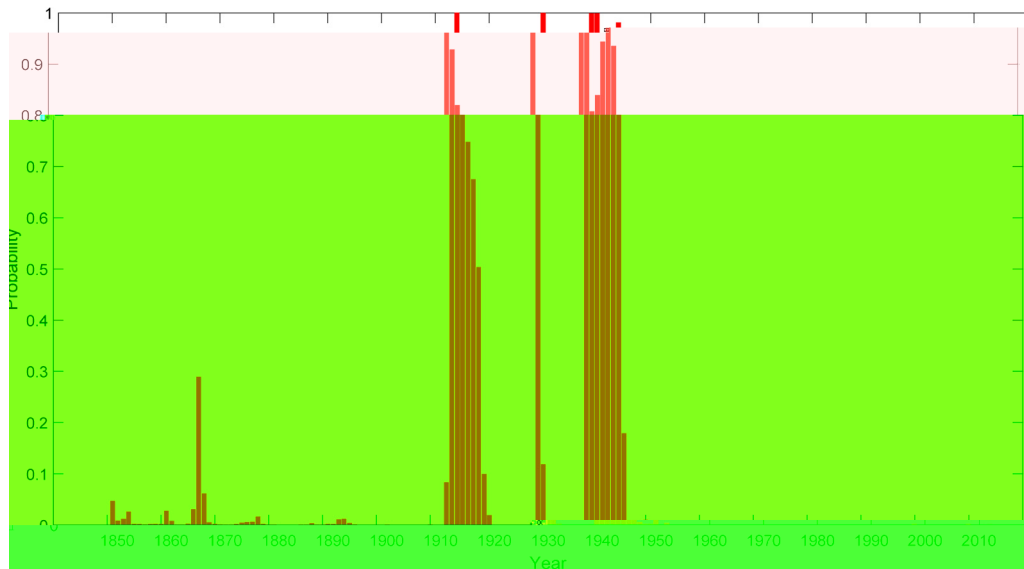
Fig. 1. World rare-event probability.
Note: This figure plots the posterior mean of the world rare-event dummy variable, I_{wt}, and, therefore, corresponds to the estimated probability that a world rare event was in effect for each year from 1851 to 2012. See equation (1) in the text.

Table 2
Country-years with Posterior Disaster Probability of 5% or More (Outside of global event years: 1867, 1914-20, 1930-31, 1939-46).

Country	Years
Argentina	1891-1902, 2001-02
Australia	1932, 1947
Belgium	1947
Brazil	1975
Canada	1921-22, 1932
Chile	1921-22, 1932-33, 1955-57, 1972-75
Colombia	1932-33, 1947-50
Denmark	1921-24, 1947-48
Egypt	1921-23, 1947-59, 1973-79
Finland	1866, 1932
Germany	1921-27, 1947-49
Greece	1947, 2009-12
Iceland	2008
India	1947-50
Malaysia	1998
Mexico	1932, 1995
New Zealand	1894-97, 1921-22, 1947-52
Norway	1921-22
Peru	1932, 1988, 1989
Portugal	1975
Russia[*]	1921-24, 1947-48
Singapore	1950-53, 1958-59
South Korea	1947-52, 1997-98
Spain	1932-38, 1947-52, 1960
Sweden	1866-69, 1921, 1947-50
Switzerland	1953-57, 1947
Taiwan	1901-12, 1947-51
Turkey	1876-81, 1887-88, 1921, 1947-50
United States	1921, 1932-33
Venezuela	1932-33, 1947-58

Note: Table 2 reports cases in which the posterior mean of the rare-event dummy variable, I_{it} for country i at time t, is at least 0.05. See equation (3) in the text.

[*] For Russia in the 1990s, the posterior disaster probability peaks at 0.14 in 1991. Using data on GDP, rather than consumption, Russia clearly shows up as a macroeconomic disaster for much of the 1990s.

Table 3
Decomposition of consumption growth.

	Mean	Share of variance of Δc_{it}	1st-order auto-correlation
Δc_{it}	0.0201	–	0.122
R_E	−0.0005	0.53	0.193
Long-run growth rate (includes LRR)	0.0023	0.10	0.876
Error term	0.0003	0.36	−0.038

Note: The entries refer to the decomposition of the annual growth rate of per capita consumption, Δc_{it}, into three parts in equation (4). R_E is the rare-events term. The term for the long-run growth rate incorporates long-run risks (LRR). The share refers to the variance in Δc_{it} associated with each term expressed as a ratio to the overall variance in Δc_{it} associated with the three terms.

crisis of 1995, the violence and economic collapse in Peru in 1985-89, the Portuguese Revolution of 1975, the Russian Revolution and civil war for 1917-24, the Spanish Civil War in 1936-38, the Korean War for South Korea for 1950-52 the Russo-Turkish War for Turkey in 1876-81, and the extended Great Depression in the United States for 1930-33.

2. Size distribution of disasters. The next group of parameters in Table 1 relates to rare events, corresponding to the R_E term in equation (4) and the dynamics of event gaps in equation (7). The parameter ρ_z determines how rapidly a country recovers from a disaster. The estimated value, 0.30 per year, implies that only 30% of a temporary disaster shock remains after one year; that is, recoveries are rapid. Note, however, that recovery refers only to the undoing of the effects from the temporary shock, $\phi_{it} - \eta_{it}$ in equation (7). The economy's consumption approaches, in the long run, a level that depends on the permanent part of the shock, η_{it}. This channel implies that there can be a great deal of long-run consequence from a disaster—depending on the realizations of η_{it} while the disaster state prevails.

The estimated mean of the disaster shock, ϕ_{it}, is −0.079; that is, consumption falls on average by about 8% in the first year of a disaster. (Note that this mean applies to a truncated normal distribution; that is, one that admits only negative values of the shock.) The estimated standard deviation, σ_ϕ, of this shock is 0.057. Hence, there is considerable dispersion in the distribution of first-year disaster sizes. The dispersion in cumulative disaster sizes depends also on the stochastic duration of disaster states, which depends on the transition probabilities given in equations (2) and (3).

The estimated mean of the permanent part of the disaster shock, η_{it}, is −0.032; that is, consumption falls on average in the long run by about 3% for each year of a disaster. (In this case, the mean applies to a normal distribution.) The estimated standard deviation, σ_η, is 0.145. Hence, there is a great deal of dispersion in the long-run consequences of a disaster.

3. LRR parameters. The final group of parameters in Table 1 concerns long-run risks (LRR), corresponding in equation (4) to the term $\chi_{i,t-1}$, which is the variable part of the long-run growth rate. The estimated value of ρ_χ, the AR(1) coefficient for χ_{it} in equation (5), is 0.73, which indicates substantial persistence from year to year. Recall that the shock to χ_{it} has a country-specific standard deviation, $k\sigma_{i,t-1}$, which evolves over time in accordance with the model of stochastic volatility in equation (6). The estimated value of ρ_σ, the AR(1) coefficient for σ_{it}^2, is 0.96, which indicates very high persistence from year to year.[7] The baseline volatility, corresponding to the mean across countries of the σ_i, is 0.024.

In key respects, our estimated parameters for the LRR part of the model accord with those presented by Bansal and Yaron (2004) and in an updated version, Bansal, Kiku, and Yaron (2010). Our estimated ρ_χ of 0.73 compares to their respective values of 0.78 and 0.74 (when their monthly values are expressed in annual terms). Our estimated ρ_σ of 0.96 compares to their respective values of 0.86 and 0.99. Our estimated mean σ_i of 0.024 compares to their respective values of 0.027 and 0.025.

From the perspective of equation (4), we can think of how the three components contribute to explaining the observed variations in the comth0.1(v)13.7(e)]TJ /F3 1 Tf 9.96 6 8 0 77.96 6 9 6797 23.8 Tm 0 Tc ( )Tj /F1 1 Tf
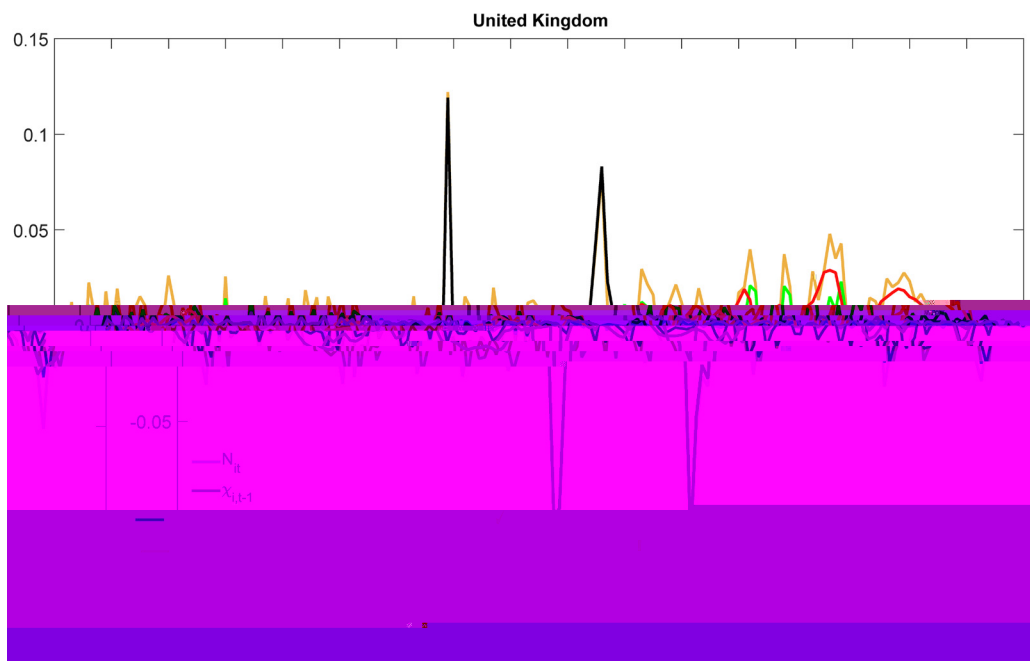
Fig. 2. Decomposition of Demeaned Consumption Growth Gap for United Kingdom.

In our present analysis, the mean recovery turns out to cumulate to 44% of the prior decline. That is, on average, 56% of the fall in consumption during a disaster is permanent. Recoveries were not considered in Barro and Ursúa (008). In Nakamura et al. (013, p. 47), the typical recovery is estimated to be 4 %.

Because the estimated standard deviation of the permanent part of the disaster shock, σ_η, is large, 0.15, there is considerable variation across disasters in the extent of recovery. In fact, simulations of the estimated model reveal that 4 percent of disasters have recoveries that exceed 100%. That is, the estimated long-run effects of many disasters are positive for the level of per capita consumption. One possible explanation is the long-term "cleansing" effects of some wars and depressions on the quality of institutions, wealth distribution, and so on. However, the estimated long-run level effect is negative in the majority of cases.

3.2. Distinctions between RE and LRR

Unlike the claim that "cyclical risks" contain disaster risks in Bansal et al. (010), the empirical results on the decomposition of growth gaps, defined in Section 2 indicate that RE and LRR are distinct risks. Figs. 2 and 3 depict the decomposition of demeaned consumption growth gaps for the United Kingdom and United States, respectively. Such figures illustrate the distinct features of the RE and LRR components. Based on the empirical identification of these components, we can summarize the rare-event component as *sporadic, drastic, and jumping outbursts* and the long-run growth rates as *persistent, moderate, and smooth fluctuations*, respectively.

The $\sigma_{i,t-1}u_{it}$ terms are essentially sequences of independent shocks, and the difference between RE_{it}^{DM} and $\sigma_{i,t-1}u_{it}$ terms are apparent. The fundamental distinctions between RE_{it}^{DM} and the long-run growth rate (or $\chi_{i,t-1}$, the persistent component of consumption growth) are as follows.

First, $\chi_{i,t-1}$ is persistent, while RE_{it}^{DM} is not. Many rare macroeconomic events burst out suddenly and unexpectedly, causing drastic changes (mostly declines) in consumption and output. Previous studies show that most of the observed macroeconomic disasters happened in periods of world disasters, such as World Wars I and II, the Great Depression, and the Great Influenza Epidemic and in periods of idiosyncratic disasters, such as regional wars, coups, and revolutions. Figs. 2 and 3 visualize the sporadic outbursts of RE_{it}^{DM}—oscillating sharply during event periods and diminishing quickly afterwards— and the persistent and smooth fluctuations of $\chi_{i,t-1}$.

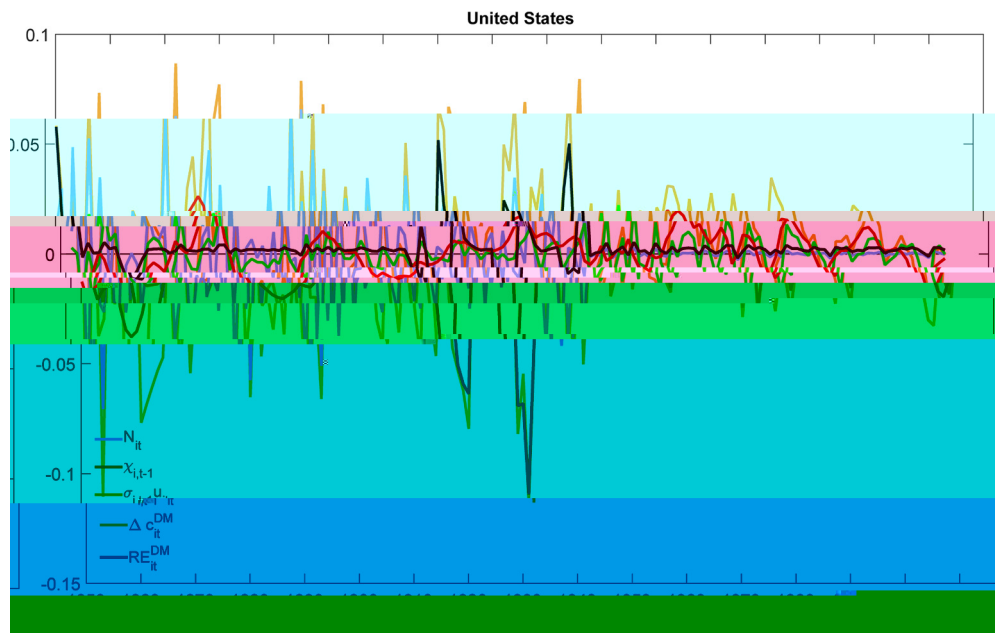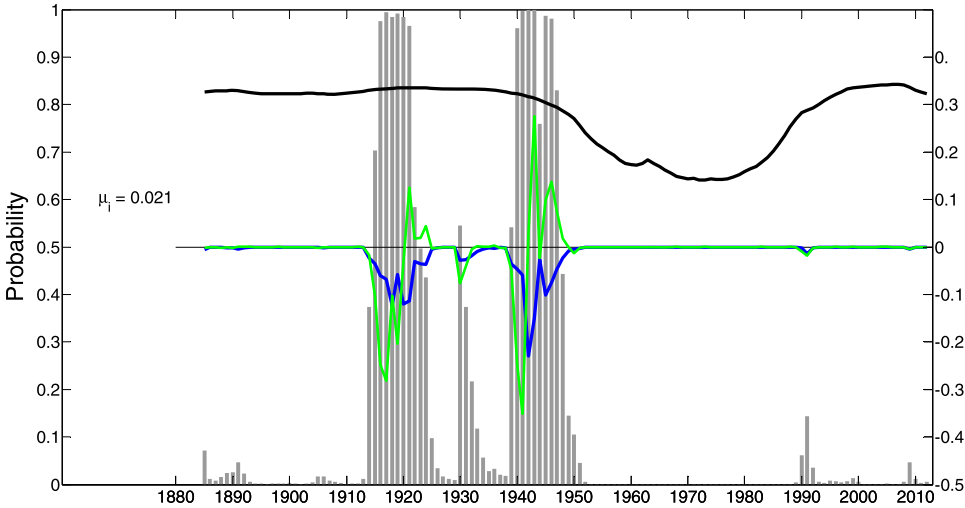Second, the volatilities of RE_{it}^{DM} and $\chi_{i,t}$

Fig. 3. Decomposition of demeaned consumption growth gap for United States.

1850 1860 1870 1880

R.J. Ba

1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0

4.2. Matching criterion: fitting the risk-free rate and return on levered equity

In this subsection, we determine the values of γ and β to fit observed long-term averages of real rates of return on corporate equity and short-term government bills (our proxy for risk-free claims). We will discuss an alternative matching criterion later. An important point here is that the parameters that describe the stochastic process for consumption were chosen solely to accord with the panel data on consumption and not to fit the data on asset returns.

For 17 countries with long-term data on asset returns, we find from an updating of Barro and Ursúa (2008, Table 5) that the average (arithmetic) real rate of return is 7.90% per year on levered equity and 0.75% per year on government bills (see Table 4, column 1). Hence, the average levered equity premium is 7.15% per year. Therefore, we calibrate the model to fit a risk-free rate of 0.75% per year and a levered equity premium of 7.15% per year. It turns out that, to fit these observations, our main analysis requires $\gamma = 5.9$ and $\beta = 0.973$.

We follow Nakamura et al. (2013) and Bansal and Yaron (2004) by making the assumption for asset pricing that the representative agent is aware contemporaneously of the values of the underlying shocks. These random variables include the indicators for a world and country-specific disaster state, the temporary and permanent shocks during disasters, the current value of the long-run growth rate, and the current level of volatility. We think that the assumption of complete current information about these underlying shocks is unrealistic. However, we also found that relaxation of this assumption had only a minor impact on the equity premium delivered by the model. The effect on the model's volatility of equity returns was more important.[9]

1. Empirical evaluation. Table 4, column 1, shows target values of various asset-pricing statistics. These targets are the mean and standard deviation of the risk-free rate, r^f, the rate of return on levered equity, r^e, and the equity premium, $r^e - r^f$; the Sharpe ratio[10]; and the mean and standard deviation of the dividend yield. These target statistics are inferred from averages in the cross-country panel data described in the notes to Table 4.

Table 4, column 2 refers to our baseline model, which combines rare events (RE) and long-run risks (LRR). Given the parameter estimates from Table 1, along with IES $= 1/\theta = 2$ (and a corporate debt-equity ratio of 0.5), the model turns out to require a coefficient of relative risk aversion, γ, of 5.9 and a subjective discount factor, β, of 0.973 (in an annual context) to fit the target values of $r^f = 0.75\%$ per year and $r^e - r^f = 7.15\%$ per year. Heuristically, we can think of γ as chosen to attain the target equity premium, with β selected to get the right overall level of rates of return.

As comparisons, Barro and Ursúa (2008) and Barro and Jin (2011) required a coefficient of relative risk aversion, γ, of 3-4 to fit the target average equity premium. In these analyses, the observed macroeconomic disasters were assumed to be fully permanent in terms of effects on the level of per capita consumption. In Nakamura et al. (2013), the required γ

Table 4
Asset-pricing statistics: data and alternative models.

Statistic	(1) Data	(2) Baseline RE & LRR	(3) RE only	(4) LRR only	(5) RE & LRR w/o stochastic volatility	(6) RE w/ perm. shocks only
mean r^f	0.0075	0.0075	0.0075	0.0075	0.0075	0.0075
mean r^e	0.0790	0.0790	0.0790	0.0790	0.0790	0.0790
mean $r^e - r^f$	0.0715	0.0715	0.0715	0.0715	0.0715	0.0715
$\sigma(r^f)$	0.0050	0.0053	0.0092	0.0131	0.0041	0.0133
$\sigma(r^e)$	0.145	0.0974	0.0861	0.0742	0.0963	0.0765
$\sigma(r^e - r^f)$	0.145	0.0872	0.0802	0.0666	0.0861	0.0698
Sharpe ratio	0.495	0.80	0.93	1.04	0.830	1.03
mean div. yield	0.0449	0.0446	0.0493	0.0457	0.0446	0.0498
σ(div. yield)	0.0175	0.0160	0.0119	0.0090	0.0147	0.0114
γ	–	5.9	6.39	17.8	5.9	6.90
β	–	0.973	0.971	0.977	0.973	0.972
mean $r^e - r^f$ with baseline parameters	–	0.0715	0.0569	0.0232	0.0665	0.0452

Notes: r^f is the risk-free rate (proxied by real returns on short-term government bills), r^e is the real total rate of return on corporate equity, σ values are standard deviations, *Sharpe ratio* is the ratio of *mean* $r^e - r^f$ to $\sigma(r^e - r^f)$, and *div. yield* is the dividend yield. A debt-equity ratio of 0.5 is assumed in the calculations for each model.

Data are means over 17 countries (Australia, Denmark, Finland, France, Germany, Italy, Japan, Netherlands, New Zealand, Norway, Spain, Sweden, Switzerland, U.K., U.S., Chile, and India) with long-term returns data, as described in Barro and Ursúa (2008, Table 5) and updated to 2014. The main underlying source is *Global Financial Data*. For the dividend yield, the means are for 8 countries with at least 90 years of data (Australia, France, Germany, Italy, Japan, Sweden, U.K., and U.S.). These data are from *Global Financial Data* and updated through 2014.

The third- and second-to-last rows give the values of γ (coefficient of relative risk aversion) and β (discount factor) required in each model to match the observed average values of the risk-free rate, r^f, and the equity return, r^e. *RE & LRR* is the baseline model, which includes all the elements of rare events (*RE*) and long-run risks (*LRR*). The other columns give results with various components eliminated. *RE only* eliminates the *LRR* parts. *LRR only* eliminates the *RE* parts. *RE & LRR, no stochastic vol.* eliminates only the stochastic volatility part of *LRR*. *RE perm. shocks only* eliminates everything except the permanent-shock part of *RE*.

The last row gives the average equity premium of each model when γ and β take on their baseline values, i.e., $\gamma = 5.9$ and $\beta = 0.973$.

of 0.5). We then recalculate for each case the values of γ and β needed to match the observed averages of 0.75% for r^f and 7.15% for $r^e - r^f$. Given these tailored parameter values, each model matches the target averages of r^f and r^e.

Table 4, column 3 (*RE only*), shows results with the omission of the long-run risks, *LRR*, parts of the model. In this case, the value of γ has to be 6.4, rather than 5.9, for the model to generate the observed average equity premium of 0.072. From this perspective, the inclusion of *LRR* in the baseline model (column 2) generates moderate improvements in the results; that is, the lower required value of γ seems more realistic. Viewed alternatively, if we retain the baseline parameter values of $\gamma = 5.9$ and $\beta = 0.973$, the model's average equity premium would fall from 0.072 (column 2) to 0.057 (column 3).

With regard to the standard deviation of r^e, the model with rare events only (column 3) has a value of 0.086, whereas the model that incorporates *LRR* has the higher value of 0.096 (column 2). In this sense, the incorporation of *LRR*

extent that the inclusion of RR improves the fit with regard to the equity premium, it is the evolution of the mean growth rate, not the fluctuation in the variance of shocks to the growth rate, that matters. With regard to the standard deviation of r^e, the value of 0.0963 in column 5 is very close to the value 0.0964 in the baseline model (column 3). In this sense, the incorporation of stochastic volatility contributes negligibly to explaining the volatility of equity returns.

Column 6 of Table 4 corresponds to using only the permanent-shock part of the rare-events, R_E model. In this case, the value of γ required to match the observed average equity premium is 6.9, not too much higher than the value 6.4 in column 3. This result shows that the main explanatory power of the R_E model for the equity premium comes from the permanent parts of rare events. Recall in this context that earlier analyses, such as Barro and Ursúa (2008) and Barro and Jin (2011), assumed that all of the rare-event shocks had fully permanent effects on the level of per capita consumption. Alternatively, if we keep the baseline parameter values of γ and β, the model's average equity premium falls from 0.057 in the full R_E model (column 3) to 0.045 (column 6). Hence, the exclusion of the temporary parts of R_E shocks has only a moderate impact on the model's average equity premium.

2. Analysis on parameter uncertainty. In the above discussion, we analyze the estimated values of γ and β when we fit the observed long-term averages of real rates of return on corporate equity and short-term government bills. A potential concern, raised by Chen et al. (2019), is that the estimated values of γ and β rely on the R_E and RR parameters, which are not known but are instead estimated from the panel data on consumption. This concern will be minor if the asset-pricing implications are robust to changes in the parameters in reasonably wide areas around the estimated values (as we later argue to be true). Chen et al. (2019) suggest that bringing in more data to identify the underlying parameters is an effective way to deal with this concern. For this reason, they think the estimations in Barro and Ursúa (2012) and Nakamura et al. (2013) work well. Therefore, it is worth noting that we utilize even more data in our present study.

To explore the robustness of the asset-pricing implications of the model, we now check the comparative statics of the asset-pricing statistics with respect to all the parameters used in our calculation. There are two sets of parameters in the calculation of asset-pricing statistics: one is the set of parameters for the consumption process, as estimated earlier, and the other set consists of the parameters for the agent's preference, namely, CRRA γ, IES $1/\theta$, and subjective discount factor β, and the debt-equity ratio ς.

As with all MCMC estimation, the estimates of the set of parameters governing the consumption process are affected by the specification of prior distributions. To avoid having the prior distributions play a biased role, we choose to make the prior distributions as "uninformative" as possible. (See the detailed discussion about priors in Appendix A.3.) The data set we are using contains 4,414 country-year observations (see Appendix A.1 for information about the data). This large macroeconomic sample helps to minimize the influence of the specification of priors.

Given the data set and priors, the MCMC estimation of parameters obeys the *square root law*: under regular conditions, statistical accuracy is inversely proportional to the square root of the Monte Carlo sample size, i.e., the length of the Markov chain used to calculate the posterior means of the parameters. According to Rosenthal (2017), under regular conditions, an MCMC asymptotic 95% confidence interval is given by $[e_n - \frac{4.4\,\hat{\sigma}_n}{\sqrt{n}}, e_n + \frac{4.4\,\hat{\sigma}_n}{\sqrt{n}}]$, where e_n is the mean estimator, $\hat{\sigma}_n$ is the standard deviation estimator, and n is the Monte Carlo sample size. In our case, for each parameter, e_n and $\hat{\sigma}_n$ are listed in Table 1, and the Monte Carlo sample size $n = 4{,}000{,}000$. As the Monte Carlo sample size is very large, the 95% confidence interval will be so narrow that the lower and upper bounds of the confidence intervals will be almost indistinguishable from the posterior means. For this reason, for each parameter, we calculate the corresponding asset-pricing statistics when the specific parameter takes on values of $e_n \pm \frac{\hat{\sigma}_n}{2}$ and other parameters are kept unchanged. The comparative statics of the asset-pricing statistics with respect to the parameters of the consumption process is shown in Table 5a and 5b.

For most of the parameters of the consumption process, the corresponding lower and upper values are relatively far apart, but the various asset-pricing statistics are close to those in the baseline model (Table 4, column 3). Generally speaking, if a change in a parameter increases the disaster risk or the long-run risk, then the model implied equity premium will be higher; otherwise, it will be lower. In Table 5a and 5b, the lowest model implied equity premium is 0.0673, which occurs when q_{10} takes on the lower value 0.694 or η takes on the upper value -0.0442. The highest model implied equity premium is 0.0769, which occurs when σ_η takes on the upper value 0.154. Note that 0.0673 and 0.0769 are only -5.9% and 7.6%, respectively, away from the baseline equity premium of 0.0715.

Table 6 shows how the results from the baseline model change with variations in the CRRA γ, subjective discount factor β, IES $1/\theta$, and debt-equity ratio ς. Column 1 has $\gamma = 4.00$, instead of the baseline value of 5.86. In other respects, the parameters are unchanged from those in Table 4, column 3. The reduction in γ lowers the model's average equity premium from 0.072 (Table 4, column 3) to 0.032 (Table 6, column 1). Conversely, Table 6, column 4, has $\gamma = 10.0$. This increase in γ raises the model's average equity premium to 0.222. Table 6, column 2 and 3 show the results for $\gamma = 5.76$ and 5.96, respectively. It is clear that the average equity premium is highly sensitive to the value of γ.

Table 6, column 5, has $\beta = 0.963$, instead of the baseline value of 0.973. The reduction in β raises r^f and r^e and lowers the equity premium. Conversely, Table 6, column 6, has $\beta = 0.983$. This increase in β lowers r^f and r^e and raises the equity premium.

Table 6, column 7, has IES $= 1/\theta = 1.5$, instead of the baseline value of 2.0. This change lowers the model's mean equity premium to 0.054. A further reduction in the IES to 1.1 (column 8) reduces the model's average equity premium further, to

For most of the parameters are

Table 5a
Asset-pricing statistics: baseline model with alternative consumption process parameters (Part I).

Parameter that deviates from the baseline model	p_0	p_1	q_{00}	q_{10}	q_{01}	q_{11}	ρ_z	$\phi^\diamond$	η
Parameter value	0.0336	0.559	0.00554	0.694	0.334	0.639	0.929	−0.0560	−0.0333
mean r^f	0.0085	0.0085	0.0077	0.0093	0.0079	0.0083	0.0075	0.0072	0.0053
mean r^e	0.0781	0.0781	0.0791	0.0771	0.0790	0.0789	0.0794	0.0796	0.0814
mean $r^e - r^f$	0.0696	0.0696	0.0714	0.0678	0.0712	0.0711	0.0719	0.0724	0.0761
$\sigma(r^f)$	0.0543	0.054	0.0552	0.0551	0.0552	0.0553	0.0555	0.0554	0.0553
$\sigma(r^e)$	0.0956	0.0960	0.0974	0.0959	0.0974	0.0975	0.0980	0.0980	0.0990
$\sigma(r^e - r^f)$	0.0885	0.0860	0.0873	0.0857	0.0873	0.0873	0.0879	0.0881	0.0888
Sharpe ratio	0.811	0.809	0.817	0.791	0.815	0.814	0.819	0.821	0.855
mean div. yield	0.0469	0.0470	0.0483	0.0464	0.0481	0.0482	0.0487	0.0488	0.0512
σ(div. yield)	0.0154	0.0156	0.0159	0.0157	0.0159	0.0159	0.0161	0.0160	0.0162

Parameter that deviates from the baseline model	p_0	p_1	q_{00}	q_{10}	q_{01}	q_{11}	ρ_z	$\phi^\diamond$	η
Parameter value	0.0346	0.792	0.00774	0.744	0.386	0.676	0.319	−0.0110	−0.0442
mean r^f	0.0066	0.0060	0.0073	0.0054	0.0072	0.0071	0.0075	0.0083	0.0096
mean r^e	0.0800	0.0812	0.0796	0.0818	0.0797	0.0798	0.0793	0.0791	0.0774
mean $r^e - r^f$	0.0734	0.0752	0.0723	0.0763	0.0726	0.0726	0.0718	0.0714	0.0678
$\sigma(r^f)$	0.0562	0.059	0.0555	0.0555	0.0554	0.0554	0.0551	0.0552	0.0553
$\sigma(r^e)$	0.0991	0.100	0.0982	0.0999	0.0981	0.0981	0.0976	0.0974	0.0967
$\sigma(r^e - r^f)$	0.0887	0.0897	0.0879	0.0897	0.0879	0.0879	0.0873	0.0872	0.0865
Sharpe ratio	0.832	0.83	0.83	0.851	0.86	0.87	0.822	0.819	0.733
mean div. yield	0.0501	0.0509	0.0490	0.0512	0.0491	0.0491	0.0486	0.0484	0.0461
σ(div. yield)	0.0165	0.0165	0.0161	0.0163	0.0161	0.0161	0.0159	0.0160	0.0158

Note: These results modify the baseline model from Table 4, column 2.

Table 5b
Asset-pricing statistics: baseline model with alternative consumption process parameters (Part II).

Parameter that deviates from the baseline model	$\sigma_\phi^\diamond$	σ_η	ρ_χ	ρ_σ	k	μ_i	σ_i^2	$\sigma_{\omega i}$	σ_{vi}
Parameter value	0.0834	0.143	0.713	0.956	0.659	0.0832	0.000472	0.0000595	0.00375
mean r^f	0.0079	0.0096	0.0079	0.0076	0.0080	0.006	0.0083	0.0079	0.0075
mean r^e	0.0790	0.0768	0.0784	0.0791	0.0782	0.0780	0.0779	0.0786	0.0793
mean $r^e - r^f$	0.0711	0.0672	0.0705	0.0715	0.0703	0.0712	0.0696	0.0707	0.0718
$\sigma(r^f)$	0.0552	0.0549	0.0552	0.0553	0.0549	0.0553	0.0549	0.0552	0.0553
$\sigma(r^e)$	0.0971	0.0962	0.0961	0.0971	0.0960	0.0974	0.0953	0.0987	0.0976
$\sigma(r^e - r^f)$	0.0886	0.0863	0.0860	0.0869	0.0862	0.0872	0.0855	0.0887	0.0874
Sharpe ratio	0.819	0.779	0.819	0.83	0.815	0.817	0.814	0.797	0.822
mean div. yield	0.0483	0.0462	0.0478	0.0485	0.0477	0.0503	0.0474	0.0483	0.0486
σ(div. yield)	0.0159	0.0156	0.0157	0.0159	0.0156	0.0161	0.0156	0.0158	0.0160

Parameter that deviates from the baseline model	$\sigma_\phi^\diamond$	σ_η	ρ_χ	ρ_σ	k	μ_i	σ_i^2	$\sigma_{\omega i}$	σ_{vi}
Parameter value	0.0954	0.154	0.747	0.970	0.752	0.0231	0.000672	0.000109	0.00655
mean r^f	0.0071	0.0052	0.0070	0.0074	0.0070	0.0082	0.0067	0.0071	0.0075
mean r^e	0.0797	0.0801	0.0805	0.0796	0.0805	0.0807	0.0809	0.0803	0.0794
mean $r^e - r^f$	0.0726	0.0769	0.0735	0.0723	0.0735	0.0725	0.0742	0.0732	0.0719
$\sigma(r^f)$	0.0555	0.0557	0.0555	0.0554	0.0557	0.0553	0.058	0.0555	0.0553
$\sigma(r^e)$	0.0985	0.0994	0.0989	0.0985	0.0996	0.0982	0.100	0.0980	0.0981
$\sigma(r^e - r^f)$	0.0885	0.0891	0.0895	0.0883	0.0891	0.0880	0.0896	0.0877	0.0880
Sharpe ratio	0.821	0.864	0.821	0.818	0.825	0.823	0.832	0.834	0.817
mean div. yield	0.0489	0.0512	0.0496	0.0489	0.0497	0.0469	0.0500	0.0495	0.0486
σ(div. yield)	0.0160	0.0164	0.0163	0.0161	0.0164	0.0159	0.0164	0.0162	0.0160

Note: These results modify the baseline model from Table 4, column 2.

0.029. Therefore, changes in the IES matter for the equity premium but, in a plausible range, not nearly as much as changes in γ.[15]

Table 6
Asset-pricing statistics: baseline model with alternative γ, β, I θ, and ς.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Parameter that deviates from the baseline model	γ	γ	γ	γ	β	β	$1/\theta$	$1/\theta$	ς	ς
Parameter value	4.00	5.76	5.96	10.0	0.963	0.83	1.50	1.10	1.00	00
mean r^f	0.0250	0.0088	0.0061	−0.0665	0.0204	−0.0055	0.0166	0.0300	0.0075	0.0075
mean r^e	0.0565	0.0775	0.0806	0.156	0.0869	0.0719	0.0708	0.0592	0.1027	0.1492
mean $r^e - r^f$	0.0315	0.0686	0.0745	0.222	0.0665	0.0773	0.0542	0.0292	0.0952	0.1417
$\sigma(r^f)$	0.0246	0.0253	0.0252	0.0222	0.0256	0.0249	0.0317	0.0434	0.0253	0.0253
$\sigma(r^e)$	0.0877	0.0969	0.0978	0.103	0.0950	0.100	0.0892	0.0761	0.115	0.173
$\sigma(r^e - r^f)$	0.0767	0.0866	0.0877	0.0893	0.0864	0.0905	0.0750	0.0600	0.104	0.170
Sharpe ratio	0.411	0.791	0.849	2.56	0.734	0.855	0.722	0.365	0.914	0.834
mean div. yield	0.0271	0.0471	0.0501	0.124	0.0564	0.0415	0.0413	0.0303	0.0625	0.0904
$\sigma(\text{div. yield})$	0.0140	0.0159	0.0160	0.0157	0.0169	0.0149	0.0171	0.0087	0.0092	0.0553

Note: These results modify the baseline model from Table 4, column 2.

Table 6, column 9, has $\varsigma = 1.0$, instead of the baseline value of 0.5. This change increases the equity premium to 0.095 and leaves r^f unchanged. A further increase in ς to 2.0 (column 10) raises the model's average equity premium further, to 0.142. Therefore, the average equity premium is sensitive to the value of ς.[16]

4.3. Alternative matching criterion

Previous studies emphasize the importance of matching the Sharpe ratio in evaluating the pricing kernel implications of economic models. (See, e.g., Hansen and Jagannathan (1991).) Thus, an alternative criterion is to match the Sharpe ratio as well as r^f and r^e. Equivalently, we can think of matching the volatility of r^e-r^f, as well as the means of r^e and r^f. A natural way to set up the matching criterion is to measure the "distance" between the model implied values and the target values of r^f, r^e, and the Sharpe ratio.

If we attach *equal importance* to matching the mean

Table 7
Asset-pricing statistics: data & various models under alternative matching criteria.

Statistic	(1) Data	(2) R E& R RR	(3) R Eonly	(4) R RR only	(5) R E& R RR w/o stochastic volatility	(6) R Ew/ perm. shocks only
mean r^f	0.0075	0.01 7	0.0136	0.01 7	0.0136	0.0137
mean r^e	0.0790	0.0340	0.0317	0.0317	0.0343	0.0 95
mean $r^e - r^f$	0.0715	0.0 14	0.0 8 2	0.0190	0.0 07	0.015 8
$\sigma(r^f)$	0.0 50	0.0 37	0.0194	0.0130	0.0 31	0.0 8 1
$\sigma(r^e)$	0. 45	0.0 31	0.06 2	0.0704	0.0 1 2	0.0613
$\sigma(r^e - r^f)$	0. 45	0.07 3	0.0614	0.064 2	0.0701	0.0534
Sharpe ratio	0. 95	0. 95	0. 96	0. 95	0. 96	0. 95
mean div. yield	0.0449	0.00 37	0.00470	0.004 4	0.00 17	0.00 26
σ(div. yield)	0.0175	0.0105	0.00 19	0.00497	0.00969	0.00793
γ	–	3.19	3 5	4.93	3.8	3.91
β	–	0.98	0.98	0.990	0.98	0.990
$\mathcal{L}(\gamma, \beta)$	–	0.0374	0.04 5	0.0411	0.0 8 4	0.046 2

Notes: For the first through the fourth-to-last rows, the data, and the setting of each model, see the notes of Table 4.
The third- and second-to-last rows give the values of $\arg\min_{(\gamma, \beta)} \mathcal{L}(\gamma, \beta)$ as in (10). The last row gives the corresponding minimum of the loss function $\mathcal{L}(\gamma, \beta)$ for each model.

It is natural to see that $\arg\min_{(\gamma, \beta)} \mathcal{L}(\gamma, \beta)$ will generate higher $\overline{r^f}$ and lower $\overline{r^e}$ than what we get in the previous subsection so as to lower the model implied Sharpe ratio S. The noticeable result is that $\arg\min_{(\gamma, \beta)} \mathcal{L}(\gamma, \beta)$ will give an almost perfect match for the Sharpe ratio, and this result is basically unchanged unless we make the denominator $\sigma^2(S)$ much larger. For instance, if we take $L(\overline{r^f}, \overline{r^e}, S)$ to be

$$L\left(\overline{r^f}, \overline{r^e}, S\right) = \frac{(\overline{r^f} - 0.0075)^2}{0.0 5^2} + \frac{(\overline{r^e} - 0.079)^2}{0. 45^2} + \frac{(S - 0. 95)^2}{0. 00^2},$$

we will have

$$\arg\min_{(\gamma, \beta)} \mathcal{L}(\gamma, \beta) = (3. 0,\ 0.9 7)$$

$$\left(\overline{r^f}, \overline{r^e}, S\right) = (0.0131,\ 0.0346, 0. 8\ ),$$

and

$$\mathcal{L}(3. 0,\ 0.9 7) = 0.0373.$$

As we can see, the model implied Sharpe ratio will now be 0. 8 , which is still very close to the target value of 0. 95. Empirical calculation shows that the estimation of (γ, β) according to criterion (10) is robust to changes in the values of the denominators $\sigma^2(r^f)$, $\sigma^2(r^e)$, and $\sigma^2(S)$.

An important point is that bringing in the Sharpe Ratio as part of the criterion for choosing the preference parameters results in a more reasonable estimate of the risk-aversion coefficient, γ, which becomes 3. 2 The downside, however, is that the model now performs poorly with respect to the equity premium, which is estimated to have a mean of only 0.0 1. From the comparison of the results for the two different matching criteria, we see there is a "trade-off" in matching the equity premium and Sharpe ratio at the same time, and it is still challenging to obtain good matches for both simultaneously.

2. Comparison of different models. Under the alternative matching criterion, the asset-pricing statistics implied by each model are shown in Table 7. The *RE & LRR* model is by far the best: It gives the smallest value of the loss function $\mathcal{L}(\gamma, \beta)$, delivers the highest equity premium of 0.0 1, and implies the lowest value of γ. The *RE & LRR w/o stochastic volatility* model is ranked second best, and the *LRR only* model performs slightly better than the *RE only* model. Note, however, that a key finding under this alternative criterion is that the highest implied equity premium (by the *RE & LRR* model) is only 0.0 1, which is much smaller than the observed value of 0.07 2

5. Time-varying disaster probability

We think that an allowance for stochastic variation in disaster probability may be an important extension to account for the remaining shortcomings in our analysis. A number of rare-disaster models argue that volatility of the disaster probability, p, or parameters that describe the size distribution of disasters is important for understanding aspects of asset pricing, notably for pricing of stock-index options. In this context, Gabaix (01) emphasizes time variation in the distribution of disaster sizes, whereas Seo and Wachter (016), Siriwardane (015), and Barro and Jiao (0 0) stress changes in disaster

probability. For most purposes, the time-varying disaster variable can be viewed as a composite of disaster probability and disaster size density.[8]

In the "normal" situation (associated with $\theta < 1$, so that the intertemporal elasticity of substitution exceeds 1), a rise in disaster probability or the typical size of a disaster lowers the price of equity. Through this channel, variations in disaster probability and sizes would impact the volatility of the rate of return on equity (and, hence, affect the Sharpe Ratio). There may also be less direct effects on means, such as the average equity premium.

The extension to allow for stochastic disaster probability is incorporated into the ongoing research of Huang et al. (2019). Due to the complexity of the numerical analysis, long-run risks have not yet been included in this analysis. In the setting where the matching criterion does not consider the Sharpe Ratio, the required coefficient of relative risk aversion γ is further reduced to 5.[2] However, the model's estimated mean Sharpe ratio is 0.675, better than previous results but still too high when compared with data. More satisfactory results in this regard will likely require the reintroduction of LRR into the model.

6. Concluding observations

Rare events (RE) and long-run risks (LRR) are complementary approaches for characterizing the long-term evolution of macroeconomic variables such as GDP and consumption. These approaches are also complementary for understanding asset-pricing patterns, including the averages of the risk-free rate and the equity premium and the volatility of equity returns. We constructed a model with RE and LRR components and estimated this joint model using long-term data on per capita consumption for 42 economies. This estimation allows us to distinguish empirically the forces associated with RE from those associated with LRR.

Rare events (RE) typically associate with major historical episodes, such as the world wars and the Great Depression and possibly the Great Influenza Pandemic (and also the ongoing coronavirus pandemic, but not the recent Great Recession). In addition to these global forces, the data reveal many disasters that affected one or a few countries. The estimated model determines the frequency and size distribution of macroeconomic disasters, including the extent and speed of eventual recovery. The distribution of recoveries is highly dispersed; that is, disasters differ greatly in terms of the relative importance of temporary and permanent components.

In contrast to RE, the long-run risks (LRR) parts of the model reflect gradual and evolving processes that apply to changing long-run growth rates and volatility. Some of these patterns relate to familiar notions about moderation and to times of persistently low or high expected growth rates.

We applied the estimated time-series model of consumption to asset pricing. A match between the model and observed average rates of return on equity and risk-free bonds requires a coefficient of relative risk aversion, γ, of 5.9. Most of the explanation for the equity premium derives from the RE components of the model, although the LRR parts make a moderate contribution. When we apply an alternative matching criterion that takes the Sharpe ratio into account, the *LRR only* model performs slightly better than the *RE only* model. Under the alternative criterion, the Sharpe ratio will be fit well, but the implied value of γ for the latter is substantially smaller and the model implied mean equity premium is very low. In other words, it is difficult to fit the equity premium and Sharpe ratio well at the same time.

We had thought that the addition of LRR to the RE framework would help to match the observed volatility of equity returns. However, the joint model still understates the volatility found in the data. Further study indicates that this aspect of the model improves if we allow for stochastic evolution of the probability or size distribution of disasters. Another extension that may further lower the required value of γ and improve the fit for the Sharpe ratio is to include a separate dividend process to which a higher leverage ratio applies.

Appendix A

A.1. Data used in this study

from data for only a few countries. The ten countries with uninterrupted data since 1851 are Denmark, France, Germany, Netherlands, Norway, Spain, Sweden, Switzerland, the United Kingdom, and the United States. The data set used in this study is much larger than those in previous studies. For example, the total number of country-year observations explored in NSBU is 1805, and that number is almost doubled here.

A.2. Missing data at the beginning of series

When $t = 1851$, i.e., for the first year in the data, the value of $I_{w,t-1}$ is missing. In this case, we use the proportion of world event years in all the years in the simulation to simulate the value of $I_{w,t-1}$ and then simulate the value of I_{wt} based on the simulated $I_{w,t-1}$ and other information.

Let t_{i0} denote the earliest date when uninterrupted consumption data are available for country i. When $t = t_{i0}$, Formula (3) is not directly applicable, because $I_{i,t_{i0}-1}$ is missing. Following the idea of (3), we calculate the following prior conditional probability instead

$$
\begin{aligned}
&\Pr(I_{it_{i0}} = 1 \mid I_{wt_{i0}}) \\
&= \Pr(I_{it_{i0}} = 1 \mid I_{i,t_{i0}-1} = 0, I_{wt_{i0}}) \Pr(I_{i,t_{i0}-1} = 0 \mid I_{wt_{i0}}) \\
&\quad + \Pr(I_{it_{i0}} = 1 \mid I_{i,t_{i0}-1} = 1, I_{wt_{i0}}) \Pr(I_{i,t_{i0}-1} = 1 \mid I_{wt_{i0}}) \\
&= q_{01}^{I_{wt_{i0}}} q_{00}^{1-I_{wt_{i0}}} \Pr(I_{i,t_{i0}-1} = 0 \mid I_{wt_{i0}}) + q_{11}^{I_{wt_{i0}}} q_{10}^{1-I_{wt_{i0}}} \Pr(I_{i,t_{i0}-1} = 1 \mid I_{wt_{i0}}).
\end{aligned}
\tag{A.1}
$$

For simplicity, we further assume

$$
\Pr(I_{i,t_{i0}-1} = 1 \mid I_{wt_{i0}}) = \Pr(I_{i,t_{i0}-1} = 1),
$$

where the prior probability $\Pr(I_{i,t_{i0}-1} = 1)$ is estimated by q_i, the fraction of event periods in all the periods studied for country i. So

$$
\Pr(I_{it_{i0}} = 1 \mid I_{wt_{i0}}) = q_{01}^{I_{wt_{i0}}} q_{00}^{1-I_{wt_{i0}}} (1 - q_i) + q_{11}^{I_{wt_{i0}}} q_{10}^{1-I_{wt_{i0}}} q_i,
\tag{A.2}
$$

and we impose the restriction that $q_i \in (0, \ 0.3]$.

For other cases of missing data, we also specify reasonable prior distributions to improve the estimation accuracy.

A.3. Prior distributions of parameters and unknown quantities

Bayesian MCMC has two major advantages in estimating the model here: (1) necessary information can be incorporated into prior beliefs, and (2) it is relatively easy to implement for a model as complicated as the one proposed in this study. The prior distributions of parameters and unknown quantities in the proposed model are listed in detail here.

In this study, a prior being "uninformative" means that the posterior distribution is proportional to the likelihood. With an uninformative prior, the mode of the posterior distribution corresponds to the maximum-likelihood estimate. A typical uninformative prior for a parameter is the uniform distribution on an infinite interval (e.g., a half-line or the entire real line). Extending that idea, we also say that the uniform distribution on a finite interval is uninformative if the finite interval contains the parameter with probability 1. More generally, we say a prior distribution is "almost uninformative" (or more rigorously, "not very informative") if it is close to a flat prior. In this study, the general guideline for the specification of priors is to make them as uninformative as possible (in certain regions). Thus, many priors are taken to be uniform.

Prior distributions of parameters   In this study, η_{it} is assumed to follow the normal distribution N

0

Non-negativity of σ_{it}^2 The method for excluding negative values of σ_{it}^2 is similar to that employed by Bansal and Yaron (2004). Instead of "replacing negative realizations with a very small number," we assume that the prior distribution of σ_{it}^2 follows the uniform distribution

$$\sigma_{it}^2 \sim U\left(10^{-8}, 0.07^2\right).$$

Thus, the posterior distribution of σ_{it}^2 follows a truncated normal distribution. This treatment is natural from the Bayesian point of view, and it is similar to that in Bansal and Yaron (2004), as both methods are using (variants of) truncated normal distributions to exclude possible negative realizations of σ_{it}^2.

A.4. Estimation procedure

The model is estimated by the Bayesian MCMC method, which has been applied to many problems in economics and finance, e.g., Chib et al. (2002); Pesaran et al. (2006); and Koop and Potter (2007). Specifically, we use the algorithm of the Gibbs sampler for the random draws of parameters and unobserved quantities (see Gelman et al. (2004) for a discussion of the MCMC algorithms).

The convergence of the MCMC simulation is guaranteed under very general conditions. In order to accurately estimate parameters and unknown quantities, we run four simulation chains, similar to the procedure in NSBU (see Appendix A.5 for details of the specification of the four simulation chains). Besides simulating multiple sequences with over-dispersed starting points throughout the parameter space and visually evaluating the trace plots of parameters and unknown quantities from the simulation, we also assess the convergence by comparing variation "between" and "within" simulated sequences (see Chapter 11 of Gelman et al. (2004) for a discussion of this method).

After a half million iterations, the simulation results from the four sets of far-apart initial values stabilize and become very close to each other. So we iterate each chain 2 million times and use the later 1 million iterations to analyze the posterior distributions of parameters and unknown quantities of interest. The first million iterations are dropped as burn-in.

A.5. Specification of four simulation chains

In order to accurately estimate the model and assess convergence, we run four independent simulation chains in a way similar to that of NSBU. We specify two extreme scenarios: one is called the "no-event scenario," the other the "all-event scenario." For the no-event scenario, we set $I_{wt} = 0$, $I_{it} = 0$, $x_{it} = c_{it}$, and $z_{it} = 0$ for all i and t. For the all-event scenario, we set $I_{wt} = 1$ and $I_{it} = 1$ for all i and t and extract a smooth trend using the Hodrick-Prescott filter (see Hodrick and Prescott (1997)). Let c_{it}^τ denote the trend component and c_{it}^c the remainder, i.e.,

$$c_{it}^c = c_{it} - c_{it}^\tau.$$

We then let

$$z_{it} = \min\left(\max\left(-0.5, c_{it}^c\right), 0\right) \quad \text{and} \quad x_{it} = c_{it} - z_{it}.$$

For each scenario, we specify two sets of initial values for parameters: one is called the "lower values," the other the "upper values." For the set of "lower values," the initial parameter values are either close to their lower bounds or very low compared to their mean values. For the "upper values," we have the opposite situation. Thus, the four sets of initial values of parameters for the four simulation chains are far apart from each other.

Appendix B. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.red.2019.08.002.

References

Balke, N.S., Gordon, R.J., 1989. The estimation of prewar Gross national product: methodology and new evidence. Journal of Political Economy 97, 38–92.
Bansal, R., Dittmar, R.F., Lundblad, C.T., 2005. Consumption, dividends, and the cross section of equity returns. The Journal of Finance 60 (4), 1639–1672.
Bansal, R., Kiku, D., Yaron, A., 2010. Long run risks, the macroeconomy, and asset prices. The American Economic Review: Papers and Proceedings 100, 542–546.
Bansal, R., Kiku, D., Yaron, A., 2016. Risks for the long run: estimation with time aggregation. Journal of Monetary Economics 82, 52–69.
Bansal, R., Shaliastovich, I., 2013. A long-run risks explanation of predictability puzzles in bond and currency markets. The Review of Financial Studies 26 (1), 1–33.
Bansal, R., Yaron, A., 2004. Risks for the long run: a potential resolution of asset pricing puzzles. The Journal of Finance 59, 1481–1509.
Barro, R.J., 2006. Rare disasters and asset markets in the twentieth century. The Quarterly Journal of Economics 121, 823–866.
Barro, R.J., 2009. Rare disasters, asset prices, and welfare costs. The American Economic Review 99, 243–264.
Barro, R.J., 2015. Environmental protection, rare disasters, and discount rates. Economica 82, 1–23.
Barro, R.J., Jin, T., 2011. On the size distribution of macroeconomic disasters. Econometrica 79, 1567–1589.
Barro, R.J., Liao, G.Y., 2019. Rare Disaster Probability and Options Pricing. Working Paper. Harvard University.

Barro, R.J., Ursúa, J.F., 2008. Macroeconomic crises since 1870. Brookings Papers on Economic Activity, 255–335.

Barro, R.J., Ursúa, J.F., 2010. Barro-Ursúa macroeconomic data. Available at https://scholar.harvard.edu/barro/publications/barro-ursua-macroeconomic-data.

Barro, R.J., Ursúa, J.F., 2012. Rare macroeconomic disasters. Annual Review of Economics 4, 83–109.

Barro, R.J., Ursúa, J.F., Weng, J., 2020. The Coronavirus and the Great Influenza Pandemic: Lessons from the 'Spanish Flu' for the Coronavirus's Potential Effects on Mortality and Economic Activity. NBER working paper 26866.

Beeler, J., Campbell, J., 2012. The long-run risks model and aggregate asset prices: an empirical assessment. Critical Finance Review 1, 141–182.

Campbell, J.Y., 2003. Consumption-based asset pricing. In: Constantinides, G., Harris, M., Stulz, R. (Eds.), Handbook of the Economics of Finance. Elsevier, Amsterdam, pp. 804–887.

Chen, H., 2010. Macroeconomic conditions and the puzzles of credit spreads and capital structure. The Journal of Finance 65, 2171–2212.

Chen, H., Dou, W.W., Kogan, L., 2019. Measuring the 'Dark Matter' in Asset Pricing Models. Working Paper. MIT.

Chib, S., Nardari, F., Shephard, N., 2002. Markov chain Monte Carlo methods for stochastic volatility models. Journal of Econometrics 108, 281–316.

Colacito, R., Croce, M.M., 2011. Risks for the long-run and the real exchange rate. Journal of Political Economy 119, 153–181.

Colacito, R., Croce, M.M., 2013. International asset pricing with recursive preferences. The Journal of Finance 68 (6), 2651–2686.

Croce, M.M., Lettau, M., Ludvigson, S.C., 2015. Investor information, long-run risk, and the term structure of equity. The Review of Financial Studies 8 (3), 706–742.

Epstein, L.G., Zin, S.E., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. Econometrica 57, 937–969.

Farhi, E., Gabaix, X., 2016. Rare disasters and exchange rates. The Quarterly Journal of Economics 131, 1–52.

Farhi, E., Fraiberger, S.P., Gabaix, X., Ranciere, R., Verdelhan, A., 2015. Crash Risk in Currency Markets. Working Paper. Harvard University.

Gabaix, X., 2012. Variable rare disasters: an exactly solved framework for ten puzzles in macro-finance. The Quarterly Journal of Economics 127, 645–700.

Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D., 2004. Bayesian Data Analysis, 2nd ed. Chapman & Hall/CRC.

Gourio, F., 2008. Disasters and recoveries. The American Economic Review 98, 68–73.

Gourio, F., 2012. Disaster risk and business cycles. The American Economic Review 102, 2734–2766.

Gruber, J., 2013. A tax-based estimate of the elasticity of intertemporal substitution. Quarterly Journal of Finance 3 (1).

Guvenen, F., 2009. A parsimonious macroeconomic model for asset pricing. Econometrica 77, 1711–1750.

Hall, R.E., 1988. Intertemporal substitution in consumption. Journal of Political Economy 96, 339–357.

Hansen, L.P., Heaton, J.C., Li, N., 2008. Consumption strikes back? Measuring long-run risk. Journal of Political Economy 116, 260–302.

Hansen, L.P., Jagannathan, R., 1991. Implications of security market data for models of dynamic economies. Journal of Political Economy 99, 225–262.

Hodrick, R.J., Prescott, E.C., 1997. Postwar U.S. business cycles: an empirical investigation. Journal of Money, Credit, and Banking 29, 1–16.

Huang, X., Jin, T., Zhou, H., 2019. (Real)istic Time-Varying Probability of Consumption Disasters. Harvard University OpenScholar Working Paper.

Koop, G., Potter, S.M., 2007. Estimation and forecasting in models with multiple breaks. The Review of Economic Studies 74, 763–789.

Malloy, C.J., Moskowitz, T.J., Vissing-Jorgensen, A., 2009. Long-run stockholder consumption risk and asset returns. The Journal of Finance 64, 2427–2479.

Mehra, R., Prescott, E.C., 1985. The equity premium: a puzzle. Journal of Monetary Economics 15, 145–161.

Nakamura, E., Steinsson, J., Barro, R.J., Ursúa, J.F., 2013. Crises and recoveries in an empirical model of consumption disasters. American Economic Journal: Macroeconomics 5, 35–74.

Nakamura, E., Sergeyev, D., Steinsson, J., 2017. Growth-rate and uncertainty shocks in consumption: cross-country evidence. American Economic Journal: Macroeconomics 9 (1), 1–39.

Pesaran, M.H., Pettenuzzo, D., Timmermann, A., 2006. Forecasting time series subject to multiple structural breaks. The Review of Economic Studies 73, 1057–1084.

Rietz, T.A., 1988. The equity risk premium: a solution. Journal of Monetary Economics 22, 117–131.

Romer, C.D., 1986. Is the stabilization of the postwar economy a figment of the data. The American Economic Review 76, 314–334.

Rosenthal, J.S., 2017. Simple confidence intervals for MCMC without CLTs. Electronic Journal of Statistics 11, 211–214.

Seo, S., Wachter, J.A., 2016. Option Prices in a Model with Stochastic Disaster Risk. Working Paper. University of Pennsylvania.

Siriwardane, E., 2015. The Probability of Rare Disasters: Estimation and Implications. Harvard Business School Finance Working Paper No. 16-061.

Wachter, J.A., 2013. Can time-varying risk of rare disasters explain aggregate stock market volatility? The Journal of Finance 68, 987–1035.

Weil, P., 1990. Nonexpected utility in macroeconomics. The Quarterly Journal of Economics 105, 29–42.