

NBER WORKING PAPER SERIES

CRISES AND RECOVERIES IN AN EMPIRICAL MODEL OF CONSUMPTION
DISASTERS

Emi Nakamura
Jón Steinsson
Robert Barro
José Ursúa

Working Paper 15920
<http://www.nber.org/papers/w15920>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2010

We would like to thank Timothy Cogley, George Constantinides, Xavier Gabaix, Ralph Koijen, Martin Lettau, Frank Schorfheide, Efthimios Tsionas, Alwyn Young, Tao Zha and seminar participants at various institutions for helpful comments and conversations. Barro would like to thank the National Science Foundation for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Emi Nakamura, Jón Steinsson, Robert Barro, and José Ursúa. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Crises and Recoveries in an Empirical Model of Consumption Disasters
Emi Nakamura, Jón Steinsson, Robert Barro, and José Ursúa
NBER Working Paper No. 15920
April 2010, Revised July 2011
JEL No. E21,G12

ABSTRACT

We estimate an empirical model of consumption disasters using a new panel data set on personal consumer expenditure for 24 countries and more than 100 years, and study its implications for asset prices. The model allows for permanent and transitory effects of disasters that unfold over multiple years. It also allows the timing of disasters to be correlated across countries. Our estimates imply that the average disaster reaches its trough after 6 years, with a peak-to-trough drop in consumption of about 30%, but that roughly half of this decline is reversed in a subsequent recovery. Uncertainty about consumption growth increases dramatically during disasters. Our estimated model generates a sizable equity premium from disaster risk, but one that is substantially smaller than in models in which disasters are permanent and instantaneous. It yields new predictions for the dynamics of risk-free interest rates, the term structure of interest rates, and the pricing of short-term versus long-term risky assets. The persistence of consumption declines in our model implies that a large value of the intertemporal elasticity of substitution is necessary to explain stock-market crashes at the onset of disasters.

Emi Nakamura
Columbia Business School
3022 Broadway, Uris Hall 820
New York, NY 10027
and NBER
enakamura@columbia.edu

Jón Steinsson
Department of Economics
Columbia University
1026 International Affairs Building
420 West 118th Street
New York, NY 10027
and NBER
jsteinsson@columbia.edu

Robert Barro
Department of Economics
Littauer Center 218
Harvard University
Cambridge, MA 02138
and NBER
rbarro@harvard.edu

José Ursúa
Department of Economics
Littauer Center G32
Harvard University
Cambridge, MA 02138
jfursua@fas.harvard.edu

1 Introduction

The average return on stocks is roughly 7% higher per year than the average return on bills across a large cross-section of countries in the twentieth century (Barro and Ursua, 2008). Mehra and Prescott (1985) argued that this large equity premium is difficult to explain in simple consumption-based asset-pricing models. A large subsequent literature in finance and macroeconomics has sought to explain this “equity-premium puzzle.” In recent years, there has been growing interest in the notion that the equity premium may be compensation for the risk of rare, but disastrous, events such as wars, depressions, and financial crises (Rietz, 1988; Barro, 2006).¹

In Barro (2006), output is a random walk with drift, and rare disasters are identified as large, instantaneous, and permanent drops in output. He calibrates the frequency and permanent impact of disasters to match large peak-to-trough drops in real per-capita GDP in a long-term panel dataset for 35 countries and shows that his model is able to match the observed equity premium with a coefficient of relative risk aversion of the representative consumer of roughly 4. More recently, Barro and Ursua (2008) have gathered a long-term data set for personal consumer expenditure in over 20 countries and shown that the same conclusions hold using these data. A growing literature has adopted this model and calibration of permanent, instantaneous disasters (e.g., Wachter, 2008; Gabaix, 2008; Farhi and Gabaix, 2008; Burnside, et al., 2008; Guo, 2007; and Gourio, 2010).²

An important critique of the Rietz-Barro disasters model calibrated to match the peak-to-trough drops in output or consumption is that it may overstate the riskiness of consumption by failing to incorporate recoveries after disasters (Gourio, 2008). A world in which disasters are followed by periods of disproportionately high growth is potentially far less risky than one in which all disasters are permanent. Kilian and Ohanian (2002) emphasize the importance of allowing for large transitory fluctuations associated with disasters such as the Great Depression and WWII in empirical models of output dynamics. More generally, a large literature in macroeconomics has debated whether it is appropriate to model output as trend or difference-stationary (Cochrane, 1988; Cogley, 1990).

A second critique of the Rietz-Barro model is that it assumes that the entire drop in output and consumption at the time of a disaster occurs instantaneously. In reality, most disasters unfold

¹Piazzesi (2010) summarizes recent research on the equity premium, emphasizing four main explanations: habits (Campbell and Cochrane, 1999), heterogeneous agents (Constantinides and Duffie, 1996), long run risk (Bansal and Yaron, 2004) and rare disasters.

²Barro and Jin (2011) show that the required coefficient of relative risk aversion can be reduced to around three if the size distribution of macroeconomic disasters is gauged by an estimated power-law distribution.

over multiple years. This profile implies that even though peak-to-trough declines in consumption exceeding 30% have occurred in many countries, the annual decline in consumption in these episodes is considerably smaller. Combining persistent declines in consumption into a single event might not be an innocuous assumption. The assumption that the entire decline in output and consumption associated with a disaster occurs in a single year is criticized in Constantinides (2008). Similarly, Julliard and Ghosh (2010) argue that using annual consumption data as opposed to peak-to-trough drops yields starkly different conclusions from Barro’s original calibration.³

Given the growing importance of the disasters model in the macroeconomics, international economics, and asset-pricing literature, a key question is whether it stands up to incorporating a more realistic process for consumption dynamics during and following disasters. We develop a model of consumption disasters that allows disasters to unfold over multiple years and to be systematically followed by recoveries. The model also allows for transitory shocks to growth in normal times and for a correlation in the timing of disasters across countries. This last feature of the model allows us to capture the fact that major disasters—such as the world wars of the 20th century—affect many countries simultaneously. Ours is the first paper to estimate the dynamic effects—both long term and short term—of these major disasters on consumption.

We estimate our model on annual consumption data from the newly constructed Barro and Ursua (2008) dataset, using Bayesian Markov-Chain Monte-Carlo (MCMC) methods.⁴ The model generates endogenous estimates of the timing, magnitude, and length of disasters, as well as the extent of recovery after disasters and the variance of shocks in disaster and non-disaster periods. Our estimation procedure also allows us to investigate the statistical uncertainty associated with the predictions of the rare-disasters model along the lines suggested by Geweke (2007) and Tsionas (2005).⁵

³Julliard and Ghosh (2010) propose a novel approach to estimating the consumption Euler equation based on generalized empirical likelihood methods, in the context of a representative agent consumption-based asset pricing model with time-additive power utility preferences. A key difference between our framework and theirs is that they focus on power utility, as in the original Rietz-Barro framework. We show that allowing for a more general preference specification is crucial in assessing the asset pricing implications of multi-period disasters and recoveries. Also, our approach does not rely on the exact timing of asset price returns during disasters. As we discuss below, asset price returns during disasters play a disproportionate role in determining the equity premium; yet these are also the periods for which asset price data are most likely to be either missing or inaccurate, for example, because of price controls during wars.

⁴We use a Metropolized Gibbs sampler. This procedure is a Gibbs sampler with a small number of Metropolis steps. See Gelfand (2000) and Smith and Gelfand (1992) for particularly lucid short descriptions of Bayesian estimation methods. See, e.g., Gelman, Carlin, Stern, and Rubin (2004) and Geweke (2005) for comprehensive treatment of these methods.

⁵In particular, we analyze the extent to which the observed asset returns are consistent with the posterior distribution of the equity premium implied by our model, taking into account parameter uncertainty. Tsionas (2005) discusses in detail the importance of accounting for finite-sample biases and parameter uncertainty in assessing the

In estimating the model, we maintain the assumption that the frequency, size distribution, and persistence of disasters is time invariant and the same for all countries. This strong assumption is important in that it allows us to pool information about disasters over time and across countries. The rare nature of disasters makes it difficult to estimate accurately a model of disasters with much variation in structural characteristics over time and space.

We find strong evidence for recoveries after disasters and for the notion that disasters unfold over several years. We estimate that disasters last roughly six years on average. Over this period, consumption drops on average by about 30% in the short run. However, about half of this drop in consumption is subsequently reversed. The average long run effect of disasters on consumption in our data is a drop of about 15%.⁶ We find that uncertainty about future consumption growth increases dramatically at the onset of a disaster. The standard deviation of consumption growth in the disaster state is roughly 12% per year, several times its value during normal times. The majority of the disasters we identify occur during World War I, the Great Depression, and World War II. Other disasters include the collapse of the Chilean economy first in the 1970's and again in the early 1980's, and the contraction in South Korea during the Asian financial crisis.

Our estimated model yields asset-pricing results that are intermediate between models that ignore disaster risk and the more parsimonious disaster models considered in the previous literature. We adopt the representative-agent endowment-economy approach to asset pricing—following Lucas (1978) and Mehra and Prescott (1985)—and assume that agents have Epstein-Zin-Weil preferences. Our model matches the observed equity premium with a coefficient of relative risk aversion (CRRA) of 6.4 and an intertemporal elasticity of substitution (IES) of 2. For these parameter values, a model without disasters yields an equity premium only one-tenth as large, while a model with one-period, permanent disasters yields an equity premium 10 times larger. Given the close link between the equity premium and the welfare costs of economic fluctuations (Alvarez and Jermann, 2004; Barro, 2009), these differences imply that our model yields costs of economic fluctuations substantially larger than a model that ignores disaster risk, but substantially smaller than the Rietz-Barro disaster model.

The differences between our model and the more parsimonious Rietz-Barro framework arise both from the recoveries and the multi-period nature of disasters. Recoveries imply that disasters

ability of alternative models to fit the observed equity premium, particularly in the presence of fat-tailed shocks.

⁶Cerra and Saxena (2008) estimate the dynamics of GDP after financial crises, civil wars and political shocks using data from 1960 to 2001 for 190 countries. They find no recovery after financial crises and political shocks but partial recovery after civil wars. Their sample does not include WWI, the Great Depression and WWII. Davis and Weinstein (2002) document a large degree of recovery at the city level after large shocks.

have a much less persistent effect on dividends, reducing the drop in stock prices when disasters occur. This modification, in turn, lowers the equity premium. The multi-period nature of disasters affects the equity premium in a more subtle way. To generate a high equity premium, the marginal utility of consumption must be high when the price of stocks drops. In our model, the price of stocks crashes at the onset of disasters—with the initial news that a disaster is underway—while consumption typically reaches its trough several years later. This lack of coincidence between the stock-market crash and the trough of consumption reduces the equity premium in our model relative to the Rietz-Barro model. In addition, since households anticipate persistent consumption declines at the onset of a disaster—they expect things to get worse before they get better—they have a strong motive to save that does not arise in the Rietz-Barro model. This desire to save limits the magnitude of the stock-market decline during disasters, further reducing the equity premium. On the other hand, if agents have EZW preferences with $CRRA > 1$ and $IES > 1$, reductions in expected future consumption growth and increases in uncertainty about future consumption raise marginal utility for a given value of current consumption.

A key feature of our model is the predictability of consumption growth during disasters—consumption typically declines for several year before recovering. These features imply that the IES, which governs consumers’ willingness to trade-off consumption over time, plays an important role in determining the asset-pricing implications of our framework. There is considerable debate in the macroeconomics and finance literature about the value of the IES. Several authors—notably Hall (1988)—argue that the IES is close to zero. However, others—such as Bansal and Yaron (2004) and Gruber (2006)—argue for substantially higher values of the IES.

The large movements in expected consumption growth associated with disasters provide a strong test of consumers’ willingness to substitute consumption over time. For a low value of the IES, our model implies a surge in stock prices at the onset of disasters and a negative equity premium in normal times. The reason is that entering the disaster state generates a strong desire to save, because consumption is expected to fall further in the short run. When the IES is substantially below one, this savings effect dominates the negative effect that the disaster has on expected future dividends from stocks and, therefore, raises the price of stocks.⁷ These predictions do not accord with the available evidence. Disasters are typically associated with stock-market crashes. This observation supports the view that consumers have a relatively high willingness to substitute

⁷Gourio (2008) makes this point forcefully in a simpler setting. For similar reasons, an IES larger than one plays an important role in the long run risk model of Bansal and Yaron (2004).

consumption over time (at least during disasters), motivating a high value of the IES.

Our estimated model yields additional predictions for the behavior of short-term and long-term interest rates. One potential concern is that the same factors driving a high equity premium would also generate a high term premium—a prediction that is not supported by the empirical evidence (Campbell, 2003; Barro and Ursua, 2008). We show that this is not the case. Our model implies a positive equity premium but a negative term premium for risk-free long-term (real) bonds that arises from the hedging properties of long-term bonds during disaster periods. Our model also generates new predictions for the dynamics of risk-free interest rates surrounding disasters. In particular, the strong desire to save during disasters drives down the return on short-term bonds, leading to low real interest rates during disaster episodes, as observed in the data.

We consider an extension of our model that allows for partial default on bonds. Empirically, inflation risk is an important source of partial default on government bonds. Data on stock and bond returns over disaster periods indicate that short-term bonds provide substantial insurance against disaster risk in only about 70% of cases. When we allow for an empirically realistic degree of default on short-term bonds, a risk aversion parameter of 7.5 is needed to fit the observed equity premium. Because inflation unfolds sluggishly in the data, the effects of inflation risk on short-term bonds is less severe than on long-term bonds. Incorporating this fact allows us to match the upward-sloping term premium for nominal bonds.

We employ the Mehra and Prescott (1985) methodology for assessing the asset-pricing implications of our model. Hansen and Singleton (1982) pioneered an alternative methodology based on measuring the empirical correlation between asset returns and the stochastic discount factor. An important difficulty with employing the Hansen-Singleton approach is that the observed timing of real returns on stocks and bonds relative to drops in consumption during disasters is affected by gaps in the data on asset prices as well as price controls, asset price controls and market closure. For example, stock price data are missing for Mexico in 1915-1918, Austria in WWII, Belgium in WWI and WWII, Portugal in 1974-1977, and Spain in 1936-1940. The Nazi regime in Germany imposed price controls in 1936 and asset-price controls in 1943 that lapsed only in 1948. In France, the stock market closed in 1940-1941 and price controls affected measured real returns over a longer period. Given these data limitations, Barro and Ursua (2009) take the approach of computing the covariance between the peak-to-trough decline in asset prices and a consumption based stochastic discount factor using a “flexible timing” assumption regarding the intervals over which these declines occur. Under this assumption, it is possible to match the equity premium for moderate

values of risk aversion. Their calculations highlight the disproportionate importance of disasters in matching the equity premium. Non-disaster periods contribute trivially to the equity premium.⁸

A number of recent papers study whether the presence of rare disasters may also help to explain other anomalous features of asset returns, such as the predictability and volatility of stock returns. These papers include Farhi and Gabaix (2008), Gabaix (2008), Gourio (2008), and Wachter (2008). Martin (2008) presents a tractable framework for asset-pricing in models of rare disasters. Gourio (2010) embeds disaster risk in a business-cycle model and shows that time-varying disaster risk can generate joint dynamics of macroeconomic aggregates and asset prices that are consistent with the data.

The paper proceeds as follows. Section 2 discusses the Barro-Ursua data on long-term personal consumer expenditure. Section 3 presents the empirical model. Section 4 discusses our estimation strategy. Section 5 presents our empirical estimates. Section 6 studies the asset-pricing implications of our model. Section 7 concludes.

2 Data

In estimating our disaster model, it is crucial to use long time series whose starting and ending points are not endogenous to the disasters themselves. It is also crucial that the data set contain information on the evolution of macroeconomic variables *during* disasters; Maddison’s (2003) tendency to interpolate GDP data during wars and other crises is not satisfactory for our purposes. Furthermore, to analyze the asset-pricing implications of rare disasters, it is important to measure consumption dynamics, as opposed to output dynamics.

We use a recently created data set on long-term personal consumer expenditures constructed by Robert Barro and Jose Ursua and described in detail in Barro and Ursua (2008).⁹ This data set includes a country only if uninterrupted annual data are available back at least before World War I, yielding a sample of 17 OECD countries and 7 non-OECD countries.¹⁰ To avoid sample-selection

⁸Another concern regarding the Hansen-Singleton methodology—emphasized by Geweke (2007) and Arakelian and Tsionas (2009)—is that parsimonious asset pricing models are sufficiently stylized that formal statistical rejections may not be very informative.

⁹These data are available from Robert Barro’s website, at:
http://www.economics.harvard.edu/faculty/barro/data_sets_barro.

¹⁰The OECD countries are: Australia, Belgium, Canada, Denmark, Finland, France, Germany, Italy, Japan, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, U.K. and U.S. The “non-OECD” countries are Argentina, Brazil, Chile, Mexico, Peru, South Korea, and Taiwan. See Barro and Ursua (2008) for a detailed description of the available data and the countries dropped due to missing data. In cases where there is a change in borders, as in the case of the unification of East and West Germany, Barro and Ursua (2008) smoothly paste together the initial per capita series for one country with that for the unified country.

bias problems associated with the starting dates of the series, we include only data after 1890. The resulting data set is an unbalanced panel of annual data for 24 countries, with data from each country starting between 1890 and 1914, yielding a total of 2685 observations.

One limitation of the Barro-Ursua consumption data set is that it does not allow us to distinguish between expenditures on non-durables and services versus durables. Unfortunately, separate data on durable and non-durable consumption are not available for most of the countries and time periods we study. For time periods when such data are available, however, the effect of excluding durables on the overall decline in consumer spending during disasters is small. The proportionate decline in spending on non durables and services is on average only 3 percentage points smaller than the overall decline in consumer spending (Barro and Ursua, 2008). The reason is that for most of the time period we study, durables accounted for only a small fraction of consumer expenditures. The effect of excluding durables is even smaller during the largest disasters, because durable consumer expenditures can at most fall to zero. The remaining fall in consumer expenditures must come entirely from non-durable expenditures.

In analyzing the asset-pricing implications of our model, we make use of total returns data on stocks, bills, and bonds from Global Financial Data (GFD), augmented with data from Dimson, Marsh, and Staunton (2002) and other sources. These data are described in detail in Barro and Ursua (2009). Unfortunately, these data are less comprehensive than the corresponding consumption series and often contain gaps for disaster periods. Price controls and controls on asset prices also make the exact timing of real returns difficult to measure during disasters. We therefore use these data to assess the predictions of our model primarily by considering average returns in non-disaster periods and cumulative returns over disaster periods.

3 An Empirical Model of Consumption Disasters

We model log consumption as the sum of three unobserved components:

$$c_{i,t} = x_{i,t} + z_{i,t} + \epsilon_{i,t}, \tag{1}$$

where $c_{i,t}$ denotes log consumption in country i at time t , $x_{i,t}$ denotes “potential” consumption in country i at time t , $z_{i,t}$ denotes the “disaster gap” of country i at time t —i.e., the amount by which consumption differs from potential due to current and past disasters—and $\epsilon_{i,t}$ denotes an i.i.d. normal shock to log consumption with a country specific variance $\sigma_{\epsilon,i,t}^2$ that potentially varies with time.

The occurrence of disasters in each country is governed by a Markov process $I_{i,t}$. Let $I_{i,t} = 0$ denote “normal times” and $I_{i,t} = 1$ denote times of disaster. The probability that a country that is not in the midst of a disaster will enter the disaster state is made up of two components: a world component and an idiosyncratic component. Let $I_{W,t}$ be an i.i.d. indicator variable that takes the value $I_{W,t} = 1$ with probability p_W . We will refer to periods in which $I_{W,t} = 1$ as periods in which “world disasters” begin. The probability that a country not in a disaster in period $t-1$ will enter the disaster state in period t is given by $p_{CbW}I_{W,t} + p_{CbI}(1 - I_{W,t})$, where p_{CbW} is the probability that a particular country will enter a disaster when a world disaster begins and p_{CbI} is the probability that a particular country will enter a disaster “on its own.” Allowing for correlation in the timing of disasters through $I_{W,t}$ is important for accurately assessing the statistical uncertainty associated with the probability of entering the disaster state. Once a country is in a disaster, the probability that it will exit the disaster state each period is p_{Ce} .

We model disasters as affecting consumption in two ways. First, disasters cause a large short-run drop in consumption. Second, disasters may affect the level of potential consumption to which the level of actual consumption will return. We model these two effects separately. First, let $\theta_{i,t}$ denote a one-off permanent shift in the level of potential consumption due to a disaster in country i at time t . Second, let $\phi_{i,t}$ denote a shock that causes a temporary drop in consumption due to the disaster in country i at time t . For simplicity, we assume that $\theta_{i,t}$ does not affect actual consumption on impact, while $\phi_{i,t}$ does not affect consumption in the long run. In this case, $\theta_{i,t}$ may represent a permanent loss of time spent on R&D and other activities that increase potential consumption or a change in institutions that the disaster induces. The short run shock, $\phi_{i,t}$, could represent destruction of structures, crowding out of consumption by government spending and temporary weakness of the financial system during the disaster.

We assume that $\theta_{i,t}$ is distributed $\theta_{i,t} \sim N(\theta, \sigma_\theta^2)$. This implies that we do not rule out the possibility that disasters can have positive long-run effects. Crises can, e.g., lead to structural change that benefits the country in the long run. We consider two distributional assumptions for the short-run shock $\phi_{i,t}$. Both of these distributions are one sided reflecting our interest in modeling disasters. In our baseline case, $\phi_{i,t}$ has a truncated normal distribution on the interval $[-\infty, 0]$. We denote this as $\phi_{i,t} \sim \text{tN}(\phi^*, \sigma_\phi^{*2}, -\infty, 0)$, where ϕ^* and σ_ϕ^{*2} denote the mean and variance, respectively, of the underlying normal distribution (before truncation). We use ϕ and σ_ϕ^2 to denote the mean and variance of the truncated distribution. We also estimate a model with $-\phi_{i,t} \sim \text{Gamma}(\alpha_\phi, \beta_\phi)$. The gamma distribution is a flexible one-sided distribution that has

excess kurtosis relative to the normal distribution.

Potential consumption evolves according to

$$\Delta x_{i,t} = \mu_{i,t} + \eta_{i,t} + I_{i,t}\theta_{i,t}, \quad (2)$$

where Δ denotes a first difference, $\mu_{i,t}$ is a country specific average growth rate of trend consumption that may vary over time, $\eta_{i,t}$ is an i.i.d. normal shock to the growth rate of trend consumption with a country specific variance $\sigma_{\eta,i}^2$. This process for potential consumption is similar to the process assumed by Barro (2006) for actual consumption. Notice that consumption in our model is trend stationary if the variances of $\eta_{i,t}$ and $\theta_{i,t}$ are zero.

The disaster gap follows an AR(1) process:

$$z_{i,t} = \rho_z z_{i,t-1} - I_{i,t}\theta_{i,t} + I_{i,t}\phi_{i,t} + \nu_{i,t}, \quad (3)$$

where $0 \leq \rho_z < 1$ denotes the first order autoregressive coefficient and $\nu_{i,t}$ is an i.i.d. normal shock with a country specific variance $\sigma_{\nu,i}^2$. We introduce $\nu_{i,t}$ mainly to aid the convergence of our numerical algorithm.¹¹ Since $\theta_{i,t}$ is assumed to affect potential consumption but to leave actual consumption unaffected on impact, it gets subtracted from the disaster gap when the disaster occurs.

Figure 1 provides an illustration of the type of disaster our model can generate. For simplicity, we abstract from trend growth and set all shocks other than $\phi_{i,t}$ and $\theta_{i,t}$ to zero. The Figure depicts a disaster that lasts six periods and in which $\rho_z = 0.6$ and $\phi_{i,t} = -0.125$ and $\theta_{i,t} = -0.0025$ in each period of the disaster. Cumulatively, log consumption drops by roughly 0.40 from peak to trough. Consumption then recovers substantially. In the long run, log consumption is 0.15 lower than it was before the disaster. This disaster is therefore partially permanent. The negative $\theta_{i,t}$ shocks during the disaster permanently lower potential consumption. The fact that the shocks to $\phi_{i,t}$ are more negative than the shocks to $\theta_{i,t}$ mean that consumption falls below potential consumption during the disaster. The difference between potential consumption and actual consumption is the disaster gap in our model. In the long run, the disaster gap closes—i.e., consumption recovers—so that only the drop in potential consumption has a long run effect on consumption. Our model can generate a wide range of paths for consumption during a disaster. If $\theta_{i,t} = 0$ throughout the

¹¹MCMC algorithms have trouble converging when the objects one is estimating are highly correlated. In our case, z_t and z_{t+j} for small j are highly correlated when there are no disturbances in the disaster gap equation between time t and time $t+j$. This would be the case in the “no disaster” periods in our model if it did not include the $\nu_{i,t}$ shock. In fact, z_t and z_{t+j} would be perfectly correlated in this case. It is in order to avoid this extremely high correlation that we introduce small disturbances to the disaster gap equation.

disaster, the entire disaster is transitory. If on the other hand $\phi_{i,t} = \theta_{i,t}$ throughout the disaster, the entire disaster is permanent.

A striking feature of the consumption data is the dramatic drop in volatility in many countries following WWII. Part of this drop in consumption volatility likely reflects changes in the procedures for constructing national accounts that were implemented at this time (Romer, 1986; Balke and Gordon, 1989). We allow for this break by assuming that $\sigma_{\epsilon,i,t}^2$ takes two values for each country: one before 1946 and one after. Allowing for this feature is important in not overestimating the occurrence of disasters in the early part of the sample. Another striking feature is that many countries experienced very rapid growth for roughly 25 years after WWII. We allow for this by assuming that $\mu_{i,t}$ takes three values for each country: one before 1946, one for the period 1946-1972 and one for the period since 1973.¹² We discuss the implications of allowing for such trend breaks in section 5.

One can show that the model is formally identified except for a few special cases in which multiple shocks have zero variance. Nevertheless, the main challenge in estimating the model is the relatively small number of disaster episodes observed in the data. We, therefore, assume that all the disaster parameters— p_W , p_{CbW} , p_{CbI} , p_{Ce} , ρ_z , θ , σ_θ^2 , ϕ , σ_ϕ^2 —are common across countries and time periods. This assumption allows us to pool information about the disasters that have occurred in different countries and at different times. In contrast, we allow the non-disaster parameters— $\mu_{i,t}$, $\sigma_{\epsilon,i,t}^2$, $\sigma_{\eta,i,t}^2$, $\sigma_{\nu,i}^2$ —to vary across countries.

4 Estimation

The model presented in section 3 decomposes consumption into three unobserved components: potential consumption, the disaster gap and a transitory shock. One way of viewing the model is, thus, as a disaster filter. Just as business-cycle filters isolate movements in output attributable to the business cycle, our model isolates movements in consumption attributable to disasters. Despite the large number of unobserved states and parameters, it is possible to estimate our model efficiently using Bayesian MCMC methods.¹³

¹²See Perron (1989) and Kilian and Ohanian (2002) for a discussion of trend breaks in macroeconomic aggregates.

¹³Bayesian MCMC methods have recently been applied to many problems in finance in which it is necessary to estimate a large number of unobserved states (see e.g., Pesaran et. al, 2006; and Koop and Potter, 2007). An important technical reason that Bayesian MCMC methods work well in our setting is that many of the unobserved states can be sampled using a Gibbs sampler as opposed to more computationally costly methods. Our algorithm samples from the posterior distributions of the parameters and states using a Gibbs sampler augmented with Metropolis steps when needed. This algorithm is described in greater detail in appendix A. The estimates discussed in section 5 for

To carry out our Bayesian estimation we need to specify a set of priors on the parameters of the model. The full set of priors we use is:

$$\begin{aligned}
\theta &\sim N(0, 0.2), & \sigma_\theta &\sim U(0.01, 0.25), \\
\phi^* &\sim U(-0.25, 0), & \sigma_\phi^* &\sim U(0.01, 0.25), \\
\phi &\sim U(-0.25, 0), & \sigma_\phi &\sim U(0.01, 0.25), \\
p_W &\sim U(0, 0.1), & p_{CbI} &\sim U(0, 0.02), \\
p_{CbW} &\sim U(0, 1), & 1 - p_{Ce} &\sim U(0, 0.9), \\
\rho_z &\sim U(0, 0.9), \\
\mu_{i,t} &\sim N(0.02, 1), & \sigma_{\epsilon,i,t} &\sim U(0, 0.15), \\
\sigma_{\eta,i} &\sim U(0, 0.15), & \sigma_{\nu,i} &\sim U(0, 0.015).
\end{aligned}$$

We consider two specifications for the short run shock $\phi_{i,t}$ —a truncated normal distribution and a gamma distribution. Thus, we specify two sets of priors for this shock. For the case of $\phi_{i,t}$ shocks that have a truncated normal distribution, we specify priors on ϕ^* and σ_ϕ^* —the mean and standard deviation of the normal distribution before it is truncated. For the alternative case with gamma distributed $\phi_{i,t}$ shocks, we place priors on the mean and standard deviation of $\phi_{i,t}$ —which we denote ϕ and σ_ϕ . These priors imply a joint prior distribution over α_ϕ and β_ϕ .

A key parameter in our model is θ —the mean long-run effect of the disaster shock, which determines the extent of recovery from a disaster. Our prior for this parameter is symmetric and highly dispersed. Thus, the prior is agnostic about whether disasters have any long run effect at all—and allows for the possibility that in some cases the long-run effect of a disaster might actually be positive, as could arise if the disaster led to a favorable change in institutions. Our estimated long run effect of disasters thus comes entirely from the data.

Our priors on the probability of disasters embed the assumption that disasters are in fact rare. On the one hand, we do not wish to “overestimate” the probability of disasters by choosing a prior on disasters that places a large prior weight on high disaster frequencies. On the other hand we do not wish to choose a prior that constrains the posterior distribution of disasters from above. In fact, our results are relatively insensitive to allowing for more dispersed priors on the probability

both versions of the model, are based on four independent Markov chains each with 2 million draws with the first 150,000 draws from each chain dropped as burn-in. The four chains are started from 2 different starting values, 2 chains from each starting value. We choose these two sets of starting values to be far apart in a sense made precise in the appendix. We use a number of techniques to assess convergence. First, we employ Gelman and Rubin’s (1992) approach to monitoring convergence based on parallel chains with “over-dispersed starting points” (see also Gelman, et al. 2004, ch 11). Second, we calculate the “effective” sample size (corrected for autocorrelation) for the parameters of the model. Finally, we visually evaluate “trace” plots from our simulated Markov chains.

of disasters, since the probability of disasters is essentially pinned down by the frequency of large and unusual events (wars, depressions, and financial crises).

Importantly, our priors in no way downweight the possibility that there are no rare disasters in the data generating process, or that the disasters are in fact small. Thus, our results on the importance of disasters are in no sense “built in” to our priors. We further verify this in section 6 by re-estimating the model using artificial data generated from a model without disasters. We show that if the model were truly generated by a process without disasters, our model would deliver a tight posterior around zero on the importance of disasters for asset prices—in stark contrast to our results based on estimating the model using actual data.

We limit the scope of disasters by setting an upper bound on the half-life of the disaster gap. This restriction rules out the possibility that consumption growth in a given period can be explained by disasters that occurred decades earlier.¹⁴ We also place upper bounds on the frequency of disasters. Our results are not sensitive to this assumption. Finally, recall that $\nu_{i,t}$ is introduced mainly to aid numerical convergence of our MCMC sampling algorithm. We therefore restrict its magnitude such that it has a negligible effect on the predictions of the model.

We have extensively investigated the robustness of our results to alternative specifications of the priors. For example, priors that restrict disasters to occur less frequently yield similar results because these specifications still allow for the infrequent occurrence of very large disasters, which contribute most to the equity premium.

5 Empirical Results

Table 1 presents our estimates of the disaster parameters for our baseline case. For each parameter, we present the parametric form of the prior distribution, the mean of the prior and its standard deviation, as well as the posterior mean and posterior standard deviation. We refer to the posterior mean of each parameter as our point estimate for that parameter.

The principle new features of our model relative to the Rietz-Barro model of permanent, instantaneous disasters are 1) the possibility of recoveries after disasters, and 2) the notion that disasters may unfold over several years. We find strong empirical support for both of these features. We can gauge the extent to which our results imply that disasters are followed by recoveries by comparing our estimate of ϕ —the mean of the short-run shock $\phi_{i,t}$ —and θ —the mean of the long-run shock $\theta_{i,t}$.

¹⁴This approach is analogous to one used in the asset-pricing literature of placing restrictions on jumps in returns and volatility (Eraker, Johannes and Polson, 2003).

We estimate $\phi = -0.111$, while we estimate $\theta = -0.025$. This implies that the short-term negative shock to consumption during disasters is on average 11.1% per year, while the long-run negative impact of the disaster on consumption is only 2.5% per year. In other words, most disasters are followed by substantial recoveries.

Our estimate of p_{Ce} —the probability that a country exits a disaster once one has begun—provides strong support for the notion that disasters unfold over several years. According to our estimates, a country that is already in a disaster will continue to be in the disaster in the following year with a 0.835 probability. This estimate implies that the average length of disasters is roughly 6 years, while the median length of disasters is 4 years.

To get a better sense for what these parameters imply about the nature of consumption disasters, Figure 2 plots the impulse response of a “typical disaster.” This prototype lasts for 6 years, and the sizes of the short-run and long-run effects are set equal to the respective posterior means of these parameters for each of the six disaster years (i.e. $\phi_{i,t} = \phi$ and $\theta_{i,t} = \theta$). The figure shows that the maximum short run effect of this typical disaster is approximately a 27% fall in consumption (a 0.32 fall in log consumption), while the long-run negative effect of the disaster is approximately 14%.¹⁵

Our estimates of σ_ϕ and σ_θ —the standard deviation of the short-run shock $\phi_{i,t}$ and long-run shock $\theta_{i,t}$ —are 0.083 and 0.121, respectively. The large estimated values of these standard deviations reveals that there is a huge amount of uncertainty during disasters about the short-run as well as the long-run effect of a disaster on consumption. Figure 3 illustrates this. Consider an agent at time 0 who knows that a disaster will begin at time 1 but knows nothing about the character of this disaster beyond the unconditional distribution. The solid line in Figure 3 plots the mean of the distribution of beliefs of such an agent about the change in log consumption going forward relative to what his beliefs were before he received the news about the start of a disaster.¹⁶ The dashed lines in the figure plot the median, 5%, and 95% quantiles of this same distribution. Figure 3 therefore gives an ex ante view of disasters, while Figure 2 gives an ex post view of a particular disaster.

Figure 3 illustrates the huge risk associated with disasters. When a disaster strikes, there is a non-trivial probability that consumption will be more than 50% lower than without the disaster

¹⁵The maximum drop is “only” roughly twice the size of the long-run drop even though the average size of the short-run shocks is more than four times larger than the average size of the long-run shock. This is because the effect of the short-run shocks in the first few years of the disaster have largely died out by the end of the disaster.

¹⁶In other words, the solid line in Figure 3 plots $E[\Delta c_{i,t+z}|I_{i,t} = 1, \Xi_{t-1}] - E[\Delta c_{i,t+z}|I_{i,t} = 0, \Xi_{t-1}]$ for $z = 0, 1, 2, \dots$ and where Ξ_{t-1} denotes the information set known to agents at time $t - 1$.

even 20-25 years later. This long left tail of the disaster distribution is particularly important for asset pricing. The median long-run effect is smaller than the mean long-run effect because the distribution of disaster sizes is negatively skewed. At first glance, Figure 3 seems to suggest more permanence in disasters than the typical disaster graph in Figure 2. This pattern arises because the average short-run effect depicted in Figure 3 averages over many disasters of varying lengths and is, therefore, muted relative to the individual disasters, which reach their troughs at different points in time.¹⁷

Figure 4 provides more detail about how our model interprets the evolution of consumption for France, Korea, Chile, and the United States.¹⁸ The two lines in each panel plot consumption and our estimate of potential consumption. The bars give our posterior probability estimate that a country was in a disaster in each year. For France, the model picks up WWI and WWII as disasters. The model views WWII as largely a transitory event for French consumption. The permanent effect of WWII on French consumption is estimated to be only about 7%. The French experience in WWII is typical for many European countries. For South Korea, our model interprets the entire period from 1940 to 1960 as a single long disaster that spans WWII and the Korean War. In contrast to the experience of many European countries, our estimates suggest that the crisis in the 1940's and 1950's had a large permanent effect on South Korean consumption (48%). This pattern is typical of the experience of Asian countries in our sample during WWII.

While the bulk of the disasters we identify are associated with world disasters, we also identify a number of idiosyncratic disaster events. Some of these idiosyncratic disasters are associated with financial or debt crises. For example, we identify a disaster in South Korea at the time of the Asian Financial Crisis and in Argentina at the time of their 2002 sovereign default.¹⁹ Other idiosyncratic disasters are associated with regional wars, coups, or revolutions. These include Chile's experience during the 1970's.

The last panel in Figure 4 plots results for the United States. Relative to most other countries in

¹⁷For example, a short disaster may reach its trough after 2 years while a long disaster may reach its trough after 10 years. The average drop in consumption at a given point in time (relative to the start of the disaster) is an average over some disaster paths for which consumption is already recovering after having reached its trough at an earlier point and other disaster paths for which consumption is still falling toward a later trough. The trough in average consumption is, therefore, far less severe than the average of the troughs across different disasters. In contrast, the long-run average level of consumption is equal to the average of the long-run levels of consumption across the different disaster paths. It is the fact that the trough in average consumption is so much less than the average of the troughs that makes the average disaster path look more permanent than in the case of the prototype disaster.

¹⁸More detailed figures for all the countries in our study are reported in a web appendix.

¹⁹Countries such as Indonesia and Thailand likely also experienced disasters during the Asian Financial Crisis but are not in the data set.

our sample, the United States was a tranquil place during our sample period. The model identifies two disaster episodes for the United States. The first disaster begins in 1914 and lasts until 1922, encompassing both WWI and the Great Influenza Epidemic of 1918-1920. The Great Depression is identified as a second disaster for U.S. consumption. The Great Depression is the larger of the two disasters with a 26% short-run drop in consumption and a 14% long-run drop.

One could also ask whether the relative tranquility of the U.S. experience since the Great Depression provides evidence that the United States is fundamentally different from other countries in our sample. However, the posterior probability for a randomly selected country experiencing no disasters over a 72-year stretch is 0.12 according to our model. The posterior probability of at least one out of 24 countries experiencing no disaster over a 73-year stretch is 0.60. Therefore, the tranquility of the U.S. experience (which is not randomly selected) does not provide evidence against our model.

Figure 5 plots our estimates of the probability that a “world disaster” began in each year.²⁰ Our model clearly identifies World War I, the Great Depression, and World War II as world disasters. Our estimate of p_W —the probability that a world disaster begins—is 3.7% per year. Countries are estimated to have a 62.3% probability of entering disasters conditional on a world disaster, but a much lower (0.6% per year) probability of entering a disaster “on their own.” The overall probability that a country enters a disaster is 2.8% per year.²¹

Our Bayesian estimation procedure does not deliver a definitive judgment on whether a disaster occurred at certain times and places but rather provides a posterior probability of whether a disaster occurred. For expositional purposes, however, it is useful to define “disaster episodes” as periods when the posterior probability of a disaster is estimated to be particularly high. We define a *disaster episode* as a set of consecutive years for a particular country such that: 1) The probability of a disaster in each of these years is larger than 10%, and 2) The sum of the probability of disaster for each year over the whole set of years is larger than one.²² In a few cases, our model is not able

²⁰This is the posterior mean of $I_{W,t}$ for each year. In other words, with the hindsight of all the data up until 2006, what is our estimate of whether a world disaster began in say 1940?

²¹The overall probability that a country will enter a disaster is $p_W p_{CbW} + (1 - p_W) p_{CbI}$. Since the three parameters involved are not independent, we cannot simply multiply together the posterior mean estimates we have for them to get a posterior mean of the overall probability of entering a disaster. Instead, we use the joint posterior distribution of these three parameters to calculate a posterior mean estimate of the overall probability that a country enters a disaster.

²²More formally: A disaster episode is a set of consecutive years for a particular country, T_i , such that for all $t \in T_i$ $P(I_{i,t} = 1) > 0.1$ and $\sum_{t \in T_i} P(I_{i,t} = 1) > 1$. The idea behind this definition is that there is a substantial posterior probability of a disaster for a particular set of consecutive years. We stress that the concept of a disaster episode is purely a descriptive device and does not influence our analysis of asset pricing. One could consider broader or narrower definitions (lower or higher cutoffs) of disaster episodes. In our experience, there are few borderline cases.

to distinguish between two or more episodes of economic turmoil that occur in the same country over a short span of time and therefore lumps these events into one long disaster episode.²³

Using this definition, we identify 53 disaster episodes. Summary statistics for the main disaster episodes are reported in Table 2, including the short-run and long-run effects of the disaster. In all cases, these statistics measure the negative effect of the disaster on the level of consumption relative to the counter-factual scenario where the country instead experienced normal trend growth. On average, the maximum drop in consumption due to the disasters is 29%, while the permanent effect of disasters on consumption is on average 14%, consistent with our estimates of the permanent and transitory components of disaster shocks.

Tables 3 and 4 present the remaining parameter estimates for our empirical model. Table 3 presents country-specific estimates of the mean growth rate of potential consumption for the countries in our sample. In most cases, the growth rate of potential consumption is estimated to have risen in 1946 and fallen in 1973, consistent with the large literature on the post-WWII “growth miracle” and the “productivity slowdown.” The structural features of the economy generating such breaks are not incorporated into our model, since investors assume that any changes in long-run growth rates they may have observed in the past will not repeat themselves in the future. An interesting question is whether there is a systematic tendency of such breaks to be positive or negative following disaster episodes. Such a pattern does not appear to be present in the data. While WWII was followed by a 30 year period of high growth in many countries, this pattern did not apply following WWI or the Great Depression. In preliminary work, Nakamura, Sergeyev, and Steinsson (2010) analyze movements in long-run growth rates around disaster periods. This framework suggests that disasters are sometimes associated with large shifts in the trend growth rate, but these changes may be positive or negative.²⁴

Table 4 presents country-specific estimates of the variances of the permanent and transitory shocks to consumption. We find a great deal of evidence for a break in the variance of the transitory shock in 1946. For all but five of the countries in our data set, our estimates of the variance of the transitory shocks to consumption fell dramatically from the earlier period to the later period. Romer (1986) argues that in the case of the United States this volatility reduction is due to improvements in measurement.

For robustness, we have estimated an alternative specification of our model in which we assume

²³Examples include WWII and the Korean war for South Korea and WWI and the Great Depression for Chile.

²⁴Bansal and Yaron’s (2004) long-run risk model suggests that persistent movements in the average growth rate of consumption and time variation in economic uncertainty could raise the equity premium implied by our model.

that $\phi_{i,t}$ —the short-run disaster shock—has a Gamma distribution. Most of the estimates are similar to the baseline case. The main difference is that the gamma model assigns a somewhat larger portion of the volatility of consumption during disasters to the short-run shock as opposed to the long-run shock.

6 Asset Pricing

We follow Mehra and Prescott (1985) in analyzing the asset-pricing implications of the consumption process we estimate in section 5 within the context of a representative consumer endowment economy. We assume that the representative consumer in our model has preferences of the type developed by Epstein and Zin (1989) and Weil (1990). For this preference specification, Epstein and Zin (1989) show that the return on an arbitrary cash flow is given by the solution to the following equation:

$$E_t \left[\beta^\xi \left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{(-\xi/\psi)} R_{w,i,t+1}^{-(1-\xi)} R_{j,i,t+1} \right] = 1, \quad (4)$$

where $R_{j,i,t+1}$ denotes the gross return on an arbitrary asset j in country i from period t to period $t+1$, and $R_{w,i,t+1}$ denotes the gross return on wealth of the representative agent in country i , which in our model equals the endowment stream. The parameter β represents the subjective discount factor of the representative consumer. The parameter $\xi = \frac{1-\gamma}{1-1/\psi}$, where γ is the coefficient of relative risk aversion (CRRA), and ψ is the intertemporal elasticity of substitution (IES), which governs the agent's desire to smooth consumption over time.²⁵

The asset-pricing implications of our model with Epstein-Zin-Weil (EZW) preferences cannot be derived analytically. We therefore use standard numerical methods.²⁶ We begin by calculating returns for two assets: a one period risk-free bill and an unleveraged claim on the consumption

²⁵The representative-consumer approach that we adopt abstracts from heterogeneity across consumers. Wilson (1968) and Constantinides (1982) show that a heterogeneous-consumer economy is isomorphic to a representative-consumer economy if markets are complete and agents have expected utility preferences. See also Rubinstein (1974). Constantinides and Duffie (1996) argue that highly persistent, heteroscedastic, uninsurable income shocks can resolve the equity-premium puzzle.

²⁶We solve the integral in equation (4) on a grid. Specifically, we start by solving for the price-dividend ratio for a consumption claim. In this case we can rewrite equation (4) as $PDR_t^C = E_t[f(\Delta C_{t+1}, PDR_{t+1}^C)]$, where PDR_t^C denotes the price dividend ratio of the consumption claim. We specify a grid for PDR_t^C over the state space. We then solve numerically for a fixed point for PDR_t^C as a function of the state of the economy on the grid. We can then rewrite equation (4) for other assets as $PDR_t = E_t[f(\Delta C_{t+1}, \Delta D_{t+1}, PDR_{t+1}^C, PDR_{t+1})]$, where PDR_t denotes the price dividend ratio of the asset in question and ΔD_{t+1} denotes the growth rate of its dividend. Given that we have already solved for PDR_t^C , we can solve numerically for a fixed point for PDR_t for any other asset as a function of the state of the economy on the grid. This approach is similar to the one used by Campbell and Cochrane (1999) and Wachter (2008).

process. In section 6.3, we calculate asset prices for a long-term bond and allow for partial default on bills and bonds during disasters.

Our asset pricing data includes rates of return for stocks, bonds and bills for 17 countries over long periods. The average arithmetic real rate of return on stocks is 8.1% per year, while the average arithmetic real rate of return on short term bills is 0.9% per year. The average equity premium is therefore 7.2% per year. If we view stock returns as a levered claim on the consumption stream, the target equity premium for an unleveraged claim on the consumption stream is lower than that for stocks. According to the Federal Reserve’s Flow-of-Funds Accounts for recent years, the debt-equity ratio for U.S. non-financial corporations is roughly one-half. This amount of leverage implies that the target equity premium for an unleveraged consumption claim in our model should be 4.8% per year $(7.2/1.5)$.²⁷ We therefore take 4.8% per year as the target for the equity premium in our analysis.

To analyze the asset-pricing implications of our model we must choose values for the CRRA, γ , the IES, ψ , and the discount factor, β . There is little agreement within the macroeconomics and finance literature about the appropriate value for the IES. Hall (1988) estimates the IES to be close to zero. This estimate is obtained by analyzing the response of aggregate consumption growth to movements in the real interest rate over time. Yet, as noted by Bansal and Yaron (2004) and Gruber (2006), the interest rate and consumption growth result from capital-market equilibrium, making it difficult to estimate the causal effect of one on the other without strong structural assumptions. These concerns are sometimes addressed by using lagged interest rates as instruments for movements in the current interest rate. However, this instrumentation strategy is successful only if there are no slowly moving parameters of preferences and technology (including especially parameters related to uncertainty) that affect interest rates and consumption growth. Alternative procedures for identifying exogenous variation in the interest rate sometimes generate much larger estimates of the IES. For example, Gruber (2006) uses instruments based on cross-state variation in tax rates on capital income to estimate a value close to 2 for the IES. As a consequence, a wide variety of parameter values for the IES are used in the asset-pricing literature. On the one hand, Campbell (2003) and Guvenen (2008) advocate values for the IES well below one, while Bansal and Yaron (2004) use a value of the IES of 1.5 and Barro (2009) relies on Gruber

²⁷Dividing the equity premium for levered equity by one plus the debt-equity ratio to get a target for unleveraged equity is exact in the simple disaster model of Barro (2006). Abel (1999) argues for approximating levered equity by a scaled consumption claim. Bansal and Yaron (2004) and others have adopted this approach. For our model, the two approaches yield virtually indistinguishable results.

(2006) to use a value of 2. We argue below that low values of the IES are starkly inconsistent with the observed behavior of asset prices during consumption disasters. We therefore focus on parameterizations with an IES equal to two— $\psi = 2$ —as our baseline case.

We present results for several different values of the CRRA. Our baseline value of the CRRA is chosen to match the equity premium in the data. Differences in the discount factor β have only minimal effects on the equity premium in our model.²⁸ They do, however, affect the risk-free rate. We choose the discount factor β to match the risk-free rate in the data for our baseline values for γ and ψ . This procedure yields a value of $\beta = \exp(-0.034)$.

The consumption data we analyze reflect any international risk sharing that agents may have engaged in. The asset-pricing equations we use are standard Euler equations involving domestic consumption and domestic asset returns. In principle, we could also investigate the asset-pricing implications of Euler equations that link domestic consumption, foreign consumption, and the exchange rate (see, e.g., Backus and Smith, 1993). A large literature in international finance explores how the form that these Euler equations take depends on the structure of international financial markets. Analyzing these issues is beyond the scope of this paper. However, recent work suggests that rare disasters may help to explain anomalies in the behavior of the real exchange rate.²⁹

6.1 The Equity Premium with Epstein-Zin-Weil Preferences

Table 6 presents our main results regarding the equity premium. The equity premium is reported for three cases: our baseline model as estimated in section 5, a version of our model without disasters as in Mehra and Prescott (1985), and a version of the model in which disasters are permanent and occur in a single period as in Barro (2006).³⁰ The statistics we report are the logarithm of the arithmetic average gross return on each asset ($\log E[R_{j,i,t+1}]$). These calculations are based on the posterior means of the parameters of our model. We discuss sampling uncertainty below.

Our estimated model matches the observed equity premium given a CRRA of 6.4. For this CRRA, the model yields an equity premium about ten times larger than the model without dis-

²⁸In the continuous time limit of our discrete time model, the equity premium is unaffected by β .

²⁹Papers on this topic include Bates (1996), Brunnermeier et al. (2008), Burnside et al. (2008), Farhi et al. (2009), Farhi and Gabaix (2008), Guo (2007) and Jurek (2008).

³⁰For the model without disasters, we set the probability of entering a disaster to zero. For the model with permanent, one-period disasters, we set the probability of exiting a disaster equal to one, assume that $\phi_{i,t} = \theta_{i,t}$, and that the distribution of these shocks corresponds to the distribution of the peak-to-trough drop in consumption over the course of disasters in our baseline model.

asters. The model without disaster risk implies essentially no equity premium, in line with Mehra and Prescott (1985). Our analysis shows, therefore, that even accounting for the partially transitory nature of disasters, and the fact that they unfold over multiple years, disaster risk greatly amplifies the equity premium. On the other hand, the model with permanent, one-period disasters of the type analyzed in Barro (2006) yields an equity premium roughly 10 times larger than our estimated model. Our analysis, thus, also shows that ignoring recoveries and the multi-year nature of disasters greatly overstates their asset-pricing implications. Given the close link between the equity premium and the welfare costs of economic fluctuations (Alvarez and Jermann, 2004; Barro, 2009), these differences imply that our model yields costs of economic fluctuations substantially larger than a model that ignores disaster risk but substantially smaller than the Rietz-Barro model of permanent and instantaneous disasters.

Figure 6 depicts equity and bond returns over the course of a “typical” disaster when $IES = 2$ and $\gamma = 6.4$. When the news arrives that a disaster has struck, the stock market crashes. In contrast, the return on risk-free bills is not affected in this initial period. This crash in the value of stocks relative to bonds at the onset of the disaster coincides with a sizable drop in consumption. The fact that stocks payoff poorly at the onset of disasters, when consumption is low and the marginal utility of consumption is high, implies that stocks must yield a considerable return-premium over bills in normal times. In other words, the equity premium in normal times in our model is compensation for the risk of a disaster occurring.

The consumption decline in any given year of a disaster is substantially smaller than the peak-to-trough declines used to calibrate simpler disaster models—we estimate the short-run effect of the disaster on consumption to be about 10% on average. In Barro (2006), disasters of a magnitude of 10% have essentially no effect on the equity premium. How, then, do our estimates generate a sizable disaster premium? The key point is that the current short-run decline in consumption is paired with news about future declines in consumption and a large increase in uncertainty about future consumption—effects that do not arise in simpler disaster models. The dramatic decline in expected future consumption growth and increase in uncertainty at the onset of the disaster contribute to the magnitude of the stock-market decline and to the premium households are willing to pay for assets that insure against disaster events.

Table 7 presents more detailed results and results for additional parameterizations. For each specification, we present results on the one hand for a long sample with a representative set of disasters and on the other hand for a long sample for which agents expect disasters to occur with

their normal frequency but no disasters actually occur. This latter case is meant to capture asset returns in “normal” times, such as the post-WWII period (at least up to 2006) in most OECD countries.

To assess the importance of allowing for recoveries after disasters, specification 2 presents asset-pricing results for the case in which disasters are completely permanent (but unfold over several years).³¹ With permanent disasters and a CRRA of 6.4, the equity premium doubles to 10%. A world in which disasters are completely permanent is clearly much riskier than a world in which there is substantial recovery after disasters. This specification of the model matches the equity premium in the data when the CRRA is set to 4.4.³² The fact that our model allows for partial recovery after disasters thus accounts for a large part of the difference in our results and the results of Barro (2006) and Barro and Ursua (2008).

To assess the role of the multi-period nature of disasters in our model, specification 3 presents results for a case in which the drop in consumption associated with a disaster occurs in a single period and the drop is permanent. With a CRRA of 6.4, this model yields an equity premium of 47%. We can match the equity premium in the data for this specification of the model with a CRRA of 3.0.³³ This specification raises the equity premium because the stock market crash coincides perfectly with the trough in consumption—when the marginal utility of consumption is highest. In contrast, when disasters unfold over multiple periods, the stock market crash occurs at the onset of the disaster, while a large fraction of the drop in consumption occurs in subsequent periods. Also, if the drop in consumption associated with a disaster occurs in a single period, it does not lead to an increased desire to save. In multi-period disasters, expectations of further drops in consumption increase the desire to save. This response strengthens the demand for stocks, limiting the magnitude of the stock-market crash.

An advantage of our formal estimation approach is that it allows us to investigate the strength of the statistical evidence for disaster risk as an explanation for the equity premium. Because they

³¹We consider a version of our model in which $\phi_{i,t} = \theta_{i,t}$ and set the mean and variance of these shocks for each year of the disaster equal to the mean and variance of peak-to-trough drops in consumption due to disasters in our baseline model divided by the expected length of disasters.

³²Notice that we lowered the CRRA by roughly 30% and this change leads to a drop in the equity premium of about 50%. This pattern illustrates that the equity premium is highly convex in the CRRA in our model. Specifications 4 and 5 of Table 7 illustrate this point further.

³³The model analyzed in specification 3 is very similar to the model analyzed by Barro and Ursua (2008). Their model matches the equity premium when $\gamma = 3.5$, while the model in specification 3 matches the equity premium for $\gamma = 3.0$. This difference arises because the size distribution of disasters in our model is relative to trend, while the peak-to-trough distribution used by Barro and Ursua (2008) does not adjust for trend growth over the course of the disaster and because of differences between our approach to estimating the distribution of disasters and the non-parametric approach used by Barro and Ursua (2008).

occur rarely, there is much less information on the frequency, size, and shape of disasters than on business-cycle phenomena. This perspective suggests that the statistical uncertainty regarding the estimates of the equity premium presented above may be large. The posterior distribution for the equity premium implied by the posterior distribution of the parameters of our model is plotted in Figure 7 for our baseline parameter values. Figure 7 shows that our estimates place more than 90% weight on parameter combinations that generate an equity premium of more than 3.3%. The centered 90% probability interval for the equity premium implied by the model is [3.0%, 7.0%].

A different way of assessing this issue is to plot the posterior distribution of the value of the CRRA that matches the observed equity premium. This distribution is plotted in Figure 8.³⁴ The centered 90% probability interval for the CRRA is [5.3, 7.8]. Thus, despite the limited data, the observed disasters provide substantial statistical evidence that it is possible to explain the observed equity premium with values of the CRRA less than 10.

To check that our results are not somehow “built in” to our priors or estimation algorithm, we analyze what our estimation algorithm implies for a data set generated from a model without disasters; that is, a setting similar to the one used by Mehra and Prescott (1985). In this counterfactual exercise, it is important that we allow ourselves only as many observations as we have in the data. We therefore simulate an artificial dataset of the same size as our data (24 countries and a total of 2685 observations) from our model with the disaster probabilities set to zero. We then estimate our model on these data and calculate the posterior distribution of the equity premium. This distribution is plotted in Figure 9. For this alternative data set, our model places a large probability on the equity premium being below 1%. These results are strikingly different from those implied by our estimated model (Figure 7), indicating that it is the data—not our priors or estimation algorithm—that lead us to the conclusion that the fear of rare disasters can explain a sizable equity premium.

It is interesting to note in Figure 9 that while the majority of the mass is located close to zero, the distribution has a long right tail. This distribution implies that even if no disasters were observed in a sample the size of ours, agents might still place some weight on the notion that disasters occur with a non-trivial probability and that the sample they had observed was simply not representative of the underlying process (a “Peso problem”).

For robustness, we also calculated asset-pricing results for the alternative specification of our

³⁴For every parameter combination sampled from the estimated posterior distributions of the parameters, we calculate the value of the CRRA required to match the equity premium.

model in which the short-run disaster shocks follow a Gamma distribution. This case yields similar asset pricing results, which are presented in specification 6 of Table 7. With $\gamma = 6.4$ and an IES of 2, the equity premium is 3.2% and the risk-free rate is 2.2%. The gamma model matches the equity premium and risk-free rate when $\gamma = 7.7$. This difference arises because the gamma model allocates slightly more of the overall volatility in consumption to the short-run shock than to the long-run shock, compared to the baseline model.

6.2 The Equity Premium with Power Utility

Much work on asset pricing—including Mehra and Prescott (1985), Rietz (1988) and Barro (2006)—considers the special case of power utility. In this case, the coefficient of relative risk aversion equals the reciprocal of the IES— $\gamma = 1/\psi$. In other words, a single parameter governs consumers’ willingness to bear risk and substitute consumption over time. Asset pricing results for our model with power utility are presented in specifications 7-9 of Table 7. With $\gamma = 1/\psi = 4$ —the utility specification used by Barro (2006)—our model yields starkly different results from those with an IES of 2. The most striking difference is that the equity premium in normal times is negative, i.e., lower than in a model in which no disasters can occur. Since the overall equity premium is positive, this model implies that high returns *during* disasters make up for low returns in normal times. This outcome contrasts with Barro (2006), in which the equity premium arises in normal times, and stocks do poorly during disasters.

Why does our model with power utility yield such different results from earlier work by Barro (2006)? The key difference is that the multi-period disasters in our model yield large movements in expected consumption growth. Figure 10 presents a time-series plot of the behavior of equity and bond returns over the course of a “typical” disaster for our baseline multi-period disaster model with power utility. Notice that there is a huge positive return on equity at the start of the disaster (when the news arrives that a disaster has struck). The reason is that entering the disaster state causes agents in the model to expect further drops in consumption going forward. Given the low value of the IES in this model (1/4) this generates a tremendous desire to save to smooth consumption that is large enough to drive up stock prices, despite the bad news about future dividends associated with the disaster. This pattern implies that agents need not be compensated for holding stocks in normal times to offset disaster risk—in fact, equity is a hedge against disaster risk and, therefore, commands a negative premium in normal times. During disasters, stockholders demand an equity premium as compensation for the risk associated with the stock-market crash that occurs at the

end of the disaster. Needless to say, the prediction that stocks yield hugely positive returns at the onset of disasters is highly counterfactual. We take this as strong evidence against low values of the IES at least during times of disaster.³⁵

These counterintuitive asset-pricing results arise because, in our estimated model, disasters unfold over multiple periods, leading to strong movements in expected consumption growth. Figure 11 presents a plot analogous to Figure 10 for the case of a single-period permanent disaster when agents have power utility. In this case, there is no change in expected consumption growth going forward, since the disaster is over as soon as it begins. As a consequence, there is no increased desire to save pushing up stock prices. Equity, thus, fares extremely poorly relative to bonds at times of disasters, and this behavior generates a large equity premium in normal times.

Another counterintuitive feature of the power utility case—emphasized by Gourio (2008)—is that one-period permanent disasters yield a lower equity premium than one-period disasters that are followed by partial recoveries—see specifications 8 and 9 in Table 7.³⁶ The reason for this difference is that, when agents expect a partial recovery after a disaster, they would like to borrow when the disaster strikes to smooth consumption. This force depresses stock prices and thus raises the equity premium. With an IES substantially below one, this force is strong enough that it outweighs the fact that the news about future dividends is not as bad in the case of partially permanent disasters as in the case of fully permanent disasters.

6.3 Long-Term Bonds, Inflation Risk, and Partial Default

The predictable movements in consumption surrounding disasters yield equilibrium movements in real interest rates that do not arise in simpler disaster models. During disasters, consumers expect consumption to keep falling and thus have an incentive to save. This force drives up the price-dividend ratio for assets and drives down their expected returns. As a consequence, stock and bill returns are lower on average during disasters than in normal times, even after the initial crash (see Figure 6). Furthermore, the return on assets is temporarily high during the recovery period after a disaster.

³⁵Similarly counter-intuitive results for the case of $IES < 1$ have been emphasized by Bansal and Yaron (2004) and Barro (2009). Bansal and Yaron (2004) observe that with an $IES < 1$ a fall in the growth rate of consumption or a rise in uncertainty leads to a rise in the price-dividend ratio of stocks. Barro (2009) shows that with an $IES < 1$ a rise in the probability of disasters also leads to a rise in the price-dividend ratio of stocks.

³⁶In specification 9, the probability of exiting a disaster equals one, implying that disasters last only one period. The distribution of $\phi_{i,t}$ is equal to the distribution of the peak-to-trough drop in consumption over the course of disasters in our baseline model. Finally, the distribution of $\theta_{i,t}$ is equal to the distribution of the long-run effect of a disaster on consumption in our baseline model.

These features of asset prices in our model line up well with the data. Barro (2006) reports low returns on bills and stocks during many disasters. He also presents evidence that real returns on U.S. Treasury bills were unusually low during wars. This regularity is inconsistent with many macroeconomic models (Barro, 1997, Ch. 12). There is, furthermore, some evidence that real returns on bills are temporarily high after wars; for example, in the United States after the Civil War and WWI. These features are absent in random-walk models of disasters, in which expected consumption growth is constant.

The variation in expected consumption growth during disasters also implies a non-trivial term structure of real interest rates. Our data contain information on real returns on long-term bonds for 15 countries over a long sample period. The underlying claims are nominal government bonds usually of around ten-year maturity. The average arithmetic real rate of return on these bonds is 2.7% per year. The real return on bills for the same sample is 1.5% per year. Thus, the average real term premium in these data is 1.2% per year.

To approximate long-term bonds in our model, we consider a perpetuity with coupon payments that decline over time. We denote the gross annual growth rate of the coupon payments by G_p . We report results for $G_p = 0.9$, a value that implies a duration for our perpetuity close to that of 10-year coupon bonds.³⁷

We begin by considering real bonds with no risk of default. The returns on such long-term bonds in our baseline model are reported in the first column of Table 8. The average return on such bonds is -2.1% per year. This result implies a term premium of -3.2% per year. In contrast, the term premium in a version of our model without disasters is virtually zero. The reason the long-term bonds have such low average returns in the presence of disasters is that they are an excellent hedge against disaster risk.

To understand why the long bond is a valuable hedge against disasters, it is useful to compare it to stocks. When a disaster occurs, stocks are affected in two ways. First, the disaster is a negative shock to future expected dividends. This effect tends to depress stock prices. Second, the representative consumer has an increased desire to save, which tends to raise stock prices. With an IES=2, the first effect dominates the second one, and stocks decline in value at the beginning of a disaster. The difference between a long-term bond and stocks is that the coupon payments on the bonds are not affected by the disaster. The only effect that the disaster has on the long-term

³⁷The duration of 10-year bonds with yields to maturity and coupon rates between 5% and 10% ranges from 6.5 years to 8 years. Our perpetuity has a duration of 7 years when its yield is 5%.

bond is therefore to raise its price because of consumers' increased desire to save. Since the price of long-term bonds rises at the onset of a disaster, these bonds provide a hedge against disaster risk and earn a lower rate of return than bills in normal times.

A potentially important feature of the data that we have so far left out of our model is the possibility of partial default on nominal bonds. While literal default on bills is rare, even during disasters, inflation may lead to partial default on bills, particularly during disasters. To calibrate the probability of partial default, we follow Barro and Ursua (2009) in considering peak-to-trough drops in stock prices over time periods that correspond roughly to consumption disasters. Extending their empirical asset-price calculations to bills, we find that in 74% of the largest consumption disasters—25 cases out of 34—stock returns are lower than bill returns.³⁸ The average stock return in these 25 cases is -34%, while the average bill return is -3%. In the remaining 9 cases, the real return on stocks and bonds are similar. In these cases, the low real returns on bills (and bonds) are caused by huge amounts of inflation. These cases also tend to be ones in which the measurement of the timing of returns is most suspect because of market closure and controls on goods and asset prices.

These calculations suggest that an appropriate calibration of the probability of partial default is 26% (9/34). To be conservative, we set the probability of partial default to 40%, as in Barro (2006). The second and third columns of Table 8 report results for calibrations that allow for partial default on bills. For a CRRA of 6.4, this modification lowers the equity premium from 4.8% to 3.3%. Raising the CRRA to 7.5 restores the equity premium to 4.8%.

The news that a disaster has struck may affect the returns on long-term bonds more than the returns on bills if it raises inflationary expectations without leading to an immediate jump in the price level. This is one possible reason for the positive term premium on long-term nominal bonds in the data. We can match this term premium by raising the probability of partial default on long-term bonds relative to short-term bonds. The fourth column in Table 8 reports results for a case in which the probability of partial default on the perpetuity is 84%, while the probability of partial default on short-term bonds is 40%. For these probabilities of partial default, our model matches the term premium on nominal bonds in the data.

Our model implies that, without default risk on bonds, the term structure is downward sloping; but introducing partial default can match the observed upward-sloping term structure for nominal

³⁸Here we identify disasters as events in which the peak-to-trough drop in consumption is larger than 17%. We choose this cutoff because applying it to the data yields a set of events that corresponds closely to the disaster episodes identified by our model. For the subset of countries that we use to estimate our model, we get 48 events as compared to 53 disaster episodes identified by our model. The average drop in consumption for these events is 32%, compared to 29% for our disaster episodes. There are 34 events for which we have data on both stock and bill returns.

bonds. If most of the default risk comes from inflation risk, our model implies that the term structure for real bonds should be less upward sloping or even downward sloping. In the United Kingdom, a large and liquid market for indexed government bonds has existed for several decades. Piazzesi and Schneider (2006) document that while the U.K. nominal yield curve has been upward sloping, the real yield curve has been downward sloping. In the United States, indexed bonds (TIPS) have been trading since 1997. Piazzesi and Schneider (2006) document that the TIPS curve over this period appears to be mostly upward sloping, contrary to our prediction. They caution, however, that this evidence is hard to assess because of the short sample and poor liquidity in the TIPS market.³⁹

6.4 Additional Predictions

Recent work by van Binsbergen, Brandt, and Koijen (2010) and Binsbergen et al. (2010) has shown that short term dividend strips on the aggregate stock market have substantially higher expected returns than the stock market as a whole.⁴⁰ They point out that this fact is difficult to match using leading equilibrium asset-pricing models, such as Campbell and Cochrane (1999), Bansal and Yaron (2004), and Barro (2006). In contrast, this feature arises naturally in our model. For the baseline calibration of our model, a one-year dividend strip on unleveraged equity—i.e., an asset that pays off a single dividend equal to consumption in the next period—has a return premium of 13.7% over bills outside of disasters, compared to an unleveraged equity premium of 4.9% for the stock market as a whole. This pattern reflects the presence of recoveries in our model, which raise the riskiness of short-term assets relative to long-term assets.⁴¹

Our model is related, in general terms, to other explanations for the equity premium in which predictable movements in consumption play an important role. A leading example is the long-run risks model of Bansal and Yaron (2004). However, one important difference between our model and theirs is that, while consumption growth is highly predictable during disasters and in the periods immediately following disasters, we do not require that it be predictable in other periods. Since

³⁹See also Evans (1998), Barr and Campbell (1997) and Campbell, Shiller, and Viceira (2009).

⁴⁰The price of a k -year dividend strip is the present value of the dividend paid in k years.

⁴¹Lettau and Wachter (2007) present a model in which long-horizon assets are less risky than short-horizon assets. A key feature of their model—like ours—is that negative shocks to dividends are associated with increases in expected dividend growth. Lettau and Wachter posit an exogenous stochastic discount factor that generates important asset pricing implications of these shocks for the cross-section of expected returns and use asset-price data to calibrate its parameters. In our model, the stochastic discount factor is derived from the utility function of the representative consumer and the dynamics of consumption. Our model therefore provides evidence of an important role for these shocks based on consumption data alone.

disasters occur infrequently, the explanatory power of the price-dividend ratio in predicting future consumption growth at medium and short horizons is close to zero in our model.⁴² In contrast, the long-run risks model generates substantial forecastability of consumption growth using the price-dividend ratio. Beeler and Campbell (2009) argue that this feature of the model is hard to reconcile with U.S. consumption data, particularly in the post-WWII period.

The analysis of this section suggests that variation in expected consumption growth surrounding disasters has important implications for asset pricing. It is nevertheless useful to ask what parameters would best approximate our estimated consumption process in the simpler random walk framework studied by Rietz (1988) and Barro (2006). The equity premium in the random-walk model is,

$$\log ER^e - \log R^f = \gamma\sigma^2 + pE\{b[(1-b)^{-\gamma} - 1]\},$$

where p denotes the probability of disasters, b denotes the permanent instantaneous fraction by which consumption drops at the time of disasters, σ^2 denotes the variance of consumption growth in normal times, and γ denotes the coefficient of relative risk aversion. Fixing the probability of disasters at our empirical estimate of 2.8%, we can replicate our baseline equity premium results in the random-walk model with a fixed disaster size of $b = 0.27$.⁴³ This value of b is substantially smaller than in the parameterizations implied by Barro (2006) and Barro and Ursua (2008) where the risk-adjusted disaster size is roughly 0.4.⁴⁴ The smaller “effective” size of disasters implied by our estimates arises from the importance of recoveries and multi-period disasters.

7 Conclusion

We estimate a quantitative model of consumption disasters that allows for recoveries, and for disasters to unfold over multiple periods. We find strong evidence for both of these features. Allowing for recoveries implies less risk associated with disasters, lowering the equity premium for given risk aversion. Allowing disasters to unfold over multiple periods implies strongly predictable movements in consumption, which also leads to a reduction in the equity premium. Even accounting

⁴²Specifically, we have analyzed regressions of consumption growth at 1, 3 and 5-year horizons on the current price dividend ratios. The R^2 of such regressions is consistently 3% or less.

⁴³Specifically, this is the value of b that matches the observed equity premium for our baseline estimate of risk aversion of $\gamma = 6.4$ and disaster probability $p = 2.8\%$.

⁴⁴Both Barro (2006) and Barro and Ursua (2008) study models with a distribution of disaster sizes. However, it is possible to solve for the value of b that matches their equity premium results for a given value of p and γ . Barro and Ursua (2008) analyze a model with a disaster probability of 0.0363 and $\gamma = 3.5$, implying that the equity premium results can be replicated with $b = 0.4$. Barro’s (2006) calibration assumes a disaster probability of 0.017 and $\gamma = 4$, implying that $b = 0.36$ is required to fit the equity premium.

for these features of the data, and for the statistical uncertainty arising from the rare nature of disasters, disaster risk greatly amplifies the equity premium.

Our estimated model matches the observed equity premium given a coefficient of relative risk aversion (CRRA) of 6.4, with a centered 90% probability interval of [5.3, 7.8]. For these parameters, a Mehra-Prescott type model that ignores disaster risk implies an equity premium close to zero. On the other hand, the Rietz-Barro model yields an equity premium more than 10 times as large as our benchmark model. These conclusions are robust to the inclusion of empirically realistic amounts of default risk on government bonds.

The predictable movements in consumption growth we estimate surrounding disasters imply that the intertemporal elasticity of substitution plays a more important role than in simpler disaster models. At the onset of a disaster, agents expect steep future declines in consumption, implying a strong desire to save. If the intertemporal elasticity of substitution is low, stock prices will counterfactually boom at the onset of disasters. This counterfactual prediction provides evidence against low values of the IES at least during times of disaster. The predictable movements in consumption we estimate also yield equilibrium movements in interest rates, a non-trivial term structure of interest rates, and predictions for dividend strips on stocks that line up well with recent empirical estimates.

An interesting extension of our approach would be to estimate a model of time-variation in disaster probabilities and trend growth rates. Aside from variation in the actual probability of disasters, the perceived disaster probability may vary due to learning. Even conditioning on all the available time-series data, our estimates suggest there is substantial uncertainty regarding the disaster parameter, implying that learning may play an important role. Variation in disaster probabilities and expected future growth rates, whether real or perceived, have the potential to generate significant volatility of asset returns—an important feature of the asset pricing data.

A Model Estimation

We employ a Bayesian MCMC algorithm to estimate our model. More specifically, we employ a Metropolized Gibbs sampling algorithm to sample from the joint posterior distribution of the unknown parameters and variables conditional on the data. This algorithm takes the following form in the case of our model.

The full probability model we employ may be denoted by

$$f(Y, X, \Theta) = f(Y, X|\Theta)f(\Theta),$$

where $Y \in \{C_{i,t}\}$ is the set of observable variables for which we have data,

$$X \in \{x_{i,t}, z_{i,t}, I_{W,t}, I_{i,t}, \phi_{i,t}, \theta_{i,t}\}$$

is the set of unobservable variables,

$$\Theta \in \{p_W, p_{CbW}, p_{CbI}, p_{Ce}, \rho_z, \theta, \sigma_\theta^2, \phi, \sigma_\phi^2, \mu_i, \sigma_{\epsilon,i,t}^2, \sigma_{\eta,i}^2, \sigma_{\nu,i}^2\}$$

is the set of parameters. From a Bayesian perspective, there is no real importance to the distinction between X and Θ . The only important distinction is between variables that are observed and those that are not. The function $f(Y, X|\Theta)$ is often referred to as the likelihood function of the model, while $f(\Theta)$ is often referred to as the prior distribution. Both $f(Y, X|\Theta)$ and $f(\Theta)$ are fully specified in sections 3 and 4 of the paper. The likelihood function may be constructed by combining equations (1)-(3), the distributional assumptions for the shocks in these equations and the distributional assumptions made about $I_{i,t}$ and $I_{W,t}$ in section 3. The prior distribution is described in detail in section 4.

The object of interest in our study is the distribution $f(X, \Theta|Y)$, i.e., the joint distribution of the unobservables conditional on the observed values of the observables. For expositional simplicity, let $\Phi = (X, \Theta)$. Using this notation, the object of interest is $f(\Phi|Y)$. The Gibbs sampler algorithm produces a sample from the joint distribution by breaking the vector of unknown variables into subsets and sampling each subvector sequentially conditional on the value of all the other unknown variables (see, e.g., Gelman et al., 2004, and Geweke, 2005). In our case we implement the Gibbs sampler as follows.

1. We derive the conditional distribution of each element of Φ conditional on all the other elements and conditional on the observables. For the i th element of Φ , we can denote this conditional distribution as $f(\Phi_i|\Phi_{-i}, Y)$, where Φ_i denotes the i th element of Φ and Φ_{-i}

denotes all but the i th element of Φ . In most cases, $f(\Phi_i|\Phi_{-i}, Y)$ are common distributions such as normal distributions or gamma distributions for which samples can be drawn in a computationally efficient manner. For example, the distribution of potential consumption for a particular country in a particular year, $x_{i,t}$, conditional on all other variables is normal. This makes using the Gibbs sampler particularly efficient in our application. Only in the case of a $(\rho_z, \sigma_{\epsilon,i,t}^2, \sigma_{\eta,i}^2, \sigma_{\nu,i}^2, \phi, \sigma_\phi^2, \sigma_\theta^2)$ are the conditional distributions not readily recognizable. In these cases, we use the Metropolis algorithm to sample values of $f(\Phi_i|\Phi_{-i}, Y)$.⁴⁵

2. We propose initial values for all the unknown variables Φ . Let Φ^0 denote these initial values.
3. We cycle through Φ sampling Φ_i^t from the distribution $f(\Phi_i|\Phi_{-i}^{t-1}, Y)$ where

$$\Phi_{-i}^{t-1} = (\Phi_1^t, \dots, \Phi_{i-1}^t, \Phi_{i+1}^{t-1}, \dots, \Phi_d^{t-1})$$

and d denotes the number of elements in Φ . At the end of each cycle, we have a new draw Φ^t . We repeat this step N times to get a sample of N draws for Φ .

4. It has been shown that samples drawn in this way converge to the distribution $f(\Phi|Y)$ under very general conditions (see, e.g., Geweke, 2005). We assess convergence and throw away an appropriate burn-in sample.

In practice, we run four such “chains” starting two from one set of initial values and two from another set of initial values. We choose starting values that are far apart in the following way: The first set of starting values has $I_{i,t} = 0$ for all i and all t and sets $x_{i,t} = c_{i,t}$ and $z_{i,t} = 0$ for all i and all t . The second set of starting values is constructed as follows. $I_{i,t} = 1$ for all i and all t . We extract a smooth trend (with breaks in 1946 and 1973) from $c_{i,t}$. Denote this trend by $c_{i,t}^t$ and denote the remaining variation in consumption as $c_{i,t}^e = c_{i,t} - c_{i,t}^t$. We set $z_{i,t} = \min(\max(-0.5, c_{i,t}^e), 0)$ and $x_{i,t} = c_{i,t} - z_{i,t}$. The first set of starting values thus attributes all the variation in the data to $x_{i,t}$, while the second attributes the bulk of the variation in the data around a smooth trend to $z_{i,t}$.

Given a sample from the joint distribution $f(\Phi|Y)$ of the unobserved variables conditional on the observed data, we can calculate any statistic of interest that involves Φ . For example, we can

⁴⁵The Metropolis algorithm samples a proposal Φ_i^* from a proposal distribution $J_t(\Phi_i^*|\Phi_i^{t-1})$. This proposal distribution must be symmetric, i.e., $J_t(x_a|x_b) = J_t(x_b|x_a)$. The proposal is accepted with probability $\min(r, 1)$ where $r = f(\Phi_i^*|\Phi_{-i}, Y)/f(\Phi_i^{t-1}|\Phi_{-i}, Y)$. If the proposal is accepted, $\Phi_i^t = \Phi_i^*$. Otherwise $\Phi_i^t = \Phi_i^{t-1}$. Using the Metropolis algorithm to sample from $f(\Phi_i|\Phi_{-i}, Y)$ is much less efficient than the standard algorithms used to sample from known distributions such as the normal distribution in most software packages. Intuitively, this is because it is difficult to come up with an efficient proposal distribution. The proposal distribution we use is a normal distribution centered at Φ_i^{t-1} .

calculate the mean of any element of Φ by calculating the sample analogue of the integral

$$\int \Phi_i f(\Phi_i | \Phi_{-i}^{t-1}, Y) d\Phi_i.$$

References

- ABEL, A. B. (1999): “Risk Premia and Term premia in General Equilibrium,” *Journal of Monetary Economics*, 43, 3–33.
- ALVAREZ, F., AND U. J. JERMANN (2004): “Using Asset Prices to Measure the Cost of Business Cycles,” *Journal of Political Economy*, 112(6), 1223–1256.
- ARAKELIAN, V., AND E. G. TSIONAS (2009): “Bayesian Analysis of the Consumption CAPM,” in *Bayesian Econometrics (Advances in Econometrics, Volume 23)*, ed. by T. B. Fomby, and R. C. Hill, pp. 619–643, England. Emerald Group Publishing.
- BACKUS, D. K., AND G. W. SMITH (1993): “Consumption and Real Exchange Rates in Dynamic Economies with Non-Traded Goods,” *Journal of International Economics*, 35(3-4), 297–316.
- BALKE, N. S., AND R. J. GORDON (1989): “The Estimation of Prewar Gross National Product: Methodology and New Evidence,” *Journal of Political Economy*, 97, 38–92.
- BANSAL, R., AND A. YARON (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 59(4), 1481–1509.
- BARR, D., AND J. Y. CAMPBELL (1997): “Inflation, Real Interest Rates, and the Bond Market: A Study of UK Nominal and Index-Linked Government Bond Prices,” *Journal of Monetary Economics*, 39, 361–383.
- BARRO, R. (1997): *Macroeconomics*. MIT Press, Cambridge, MA.
- (2006): “Rare Disasters and Asset Markets in the Twentieth Century,” *Quarterly Journal of Economics*, 121(3), 832–866.
- (2009): “Rare Disasters, Asset Prices, and Welfare Costs,” *American Economic Review*, 99(1), 243–264.
- BARRO, R. J., AND T. JIN (2011): “On the Size Distribution of Macroeconomic Disasters,” *Econometrica*, forthcoming.
- BARRO, R. J., AND J. F. URSUA (2008): “Macroeconomic Crises since 1870,” *Brookings Papers on Economic Activity*, 2008, 255–350.
- (2009): “Stock-Market Crashes and Depressions,” NBER Working Paper No. 14760.
- BATES, D. (1996): “Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options,” *Journal of Financial Economics*, 9, 69–107.
- BEELER, J., AND J. Y. CAMPBELL (2009): “The Long-Run Risk Model and Aggregate Asset Prices: An Empirical Assessment,” NBER Working Paper No. 14788.

- BRUNNERMEIER, M. K., S. NAGEL, AND L. H. PEDERSEN (2008): "Carry Trades and Currency Crashes," in *NBER Macroeconomics Annual*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford, pp. 313–347, Chicago, IL: University of Chicago Press.
- BURNSIDE, A., M. S. EICHENBAUM, I. KLESCHCHESKI, AND S. REBELO (2008): "Do Peso Problems Explain the Returns to the Carry Trade?," NBER Working Paper No. 14054.
- CAMPBELL, J. Y. (2003): "Consumption-Based Asset Pricing," in *Handbook of the Economics of Finance*, ed. by G. Constantinides, M. Harris, and R. Stulz, pp. 804–887, Amsterdam: Elsevier.
- CAMPBELL, J. Y., AND J. H. COCHRANE (1999): "By force of habit: A consumption-based explanation of aggregate stock market behavior," *Journal of Political Economy*, 107, 205–251.
- CAMPBELL, J. Y., R. SHILLER, AND L. VICEIRA (2009): "Understanding Inflation-Indexed Bond Markets," *Brookings Papers on Economic Activity*, 2009(1), 79–120.
- CERRA, V., AND S. C. SAXENA (2008): "Growth Dynamics: The Myth of Economic Recovery," *American Economic Review*, 98(1), 439–457.
- COCHRANE, J. H. (1988): "How Big Is the Random Walk in GNP?," *Journal of Political Economy*, 96(5), 893–920.
- COGLEY, T. (1990): "International Evidence on the Size of the Random Walk in Output," *Journal of Political Economy*, 98(3), 501–518.
- CONSTANTINIDES, G. (1982): "Intertemporal Asset Pricing with Heterogeneous Consumers and without Demand Aggregation," *Journal of Business*, 55(2), 253–267.
- (2008): "Consumption Crises since 1870," Brookings Conference Discussion of "Macroeconomic Crises since 1870" by Robert J. Barro and Jose F. Ursua.
- CONSTANTINIDES, G., AND D. DUFFIE (1996): "Asset Pricing with heterogeneous consumers," *Journal of Political Economy*, 104(2), 219–240.
- DAVIS, D. R., AND D. E. WEINSTEIN (2002): "Bones, Bombs and Break Points: The Geography of Economic Activity," *American Economic Review*, 92(5), 1269–1289.
- DIMSON, E., P. MARSH, AND M. STAUNTON (2002): *Triumph of the Optimists*. Princeton University Press, Princeton, New Jersey.
- EPSTEIN, L. G., AND S. E. ZIN (1989): "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns," *Econometrica*, 57, 937–969.
- ERAKER, B., M. JOHANNES, AND N. POLSON (2003): "The Impact of Jumps in Volatility and Returns," *Journal of Finance*, 58, 1269–1300.
- EVANS, M. (1998): "Real Rates, Expected Inflation, and Inflation Risk Premia," *Journal of Finance*, 53(1), 187–218.
- FARHI, E., S. P. FRAIBERGER, X. GABAIX, R. RANCIERE, AND A. VERDELHAN (2009): "Crash Risk in Currency Markets," Working Paper, Harvard University.
- FARHI, E., AND X. GABAIX (2008): "Rare Disasters and Exchange Rates," Working Paper.
- GABAIX, X. (2008): "Variable Rare Disasters: A Tractable Theory of Ten Puzzles in Macro-Finance," *American Economic Review*, 98(2), 64–67.

- GELFAND, A. (2000): “Gibbs Sampling,” *Journal of the American Statistical Association*, 95(452), 1300–1304.
- GELMAN, A., J. B. CARLIN, H. S. STERN, AND D. B. RUBIN (2004): *Bayesian Data Analysis*. John Wiley and Sons, Hoboken, New Jersey.
- GELMAN, A., AND D. B. RUBIN (1992): “Inference from Iterative Simulation Using Multiple Sequences,” *Statistical Science*, 7, 457–511.
- GEWEKE, J. (2005): *Contemporary Bayesian Econometrics and Statistics*. Chapman & Hall/CRC, Boca Raton, Florida.
- (2007): “Models, Computational Experiments and Reality,” Working Paper, Iowa University.
- GOURIO, F. (2008): “Disasters and Recoveries,” *American Economic Review*, 98(2), 68–73.
- (2010): “Disaster Risk and Business Cycles,” Working Paper, Boston University.
- GRUBER, J. (2006): “A Tax-Based Estimate of the Elasticity of Intertemporal Substitution,” NBER Working Paper No. 11945.
- GUO, K. (2007): “Exchange Rates and Asset Prices in an Open Economy with Rare Disasters,” Working Paper, Harvard University.
- GUVENEN, F. (2008): “A Parsimonious Macroeconomic Model for Asset Pricing,” Working Paper.
- HALL, R. E. (1988): “Intertemporal Substitution in Consumption,” *Journal of Political Economy*, 96(2), 339–357.
- HANSEN, L. P., AND K. J. SINGLETON (1982): “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models,” *Econometrica*, 50, 1269–1288.
- JULLIARD, C., AND A. GHOSH (2010): “Can Rare Events Explain the Equity Premium Puzzle?,” Working Paper.
- JUREK, J. W. (2008): “Crash-Neutral Currency Carry Trades,” Working Paper, Princeton University.
- KILIAN, L., AND L. E. OHANIAN (2002): “Unit Roots, Trend Breaks, and Transitory Dynamics: A Macroeconomic Perspective,” *Macroeconomic Dynamics*, 6(5), 614–632.
- KOOP, G., AND S. M. POTTER (2007): “Estimation and Forecasting in Models with Multiple Breaks,” *Review of Economic Studies*, 74(3), 763–789.
- LETTAU, M., AND J. A. WACHTER (2007): “Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium,” *Journal of Finance*, 62(1), 55–92.
- LUCAS, R. E. (1978): “Asset Prices in an Exchange Economy,” *Econometrica*, 46(6), 1429–1445.
- MADDISON, A. (2003): *The World Economy: Historical Statistics*. OECD, Paris.
- MARTIN, I. (2008): “Disasters and the Welfare Cost of Uncertainty,” *American Economic Review*, 98(2), 74–78.
- MEHRA, R., AND E. C. PRESCOTT (1985): “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 15, 145–161.

- NAKAMURA, E., D. SERGEYEV, AND J. STEINSSON (2010): “Long Run Risk and the Equity Premium: Evidence from Long Term Consumption Data,” Working Paper, Columbia University.
- PERRON, P. (1989): “The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis,” *Econometrica*, 57, 1361–1401.
- PESARAN, M. H., D. PETTENUZZO, AND A. TIMMERMAN (2006): “Forecasting Time Series Subject to Multiple Structural Breaks,” *Review of Economic Studies*, 73(4), 1057–1084.
- PIAZZESI, M. (2010): “Understanding Returns on Average and over Time,” *NBER Reporter 2010*, (2), Program Report for the NBER Asset Pricing Program.
- PIAZZESI, M., AND M. SCHNEIDER (2006): “Equilibrium Yield Curves,” in *NBER Macroeconomics Annual*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford, pp. 389–442, Cambridge, Ma. MIT Press.
- 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,” *American Economic Review*, 76(3), 314–334.
- RUBINSTEIN, M. (1974): “An Aggregation Theorem for Securities Markets,” *Journal of Financial Economics*, 1, 225–244.
- SMITH, A., AND A. GELFAND (1992): “Bayesian Statistics without Tears: A Sampling-Resampling Perspective,” *The American Statistician*, 46(2), 84–88.
- TSIONAS, E. G. (2005): “Likelihood Evidence on the Asset Returns Puzzle,” *Review of Economic Studies*, 72, 917–946.
- VAN BINSBERGEN, J. H., M. W. BRANDT, AND R. S. KOIJEN (2010): “On the Timing and Pricing of Dividends,” Working Paper, Northwestern University.
- VAN BINSBERGEN, J. H., W. HUESKES, R. S. KOIJEN, AND E. VRUGT (2010): “Asset Pricing Puzzles: Measuring the Importance of Short-Term and Long-Term Risk,” Working Paper, Northwestern University.
- WACHTER, J. A. (2008): “Can time-varying risk of rare disasters explain aggregate stock market volatility,” Working Paper.
- WEIL, P. (1990): “Nonexpected Utility in Macroeconomics,” *Quarterly Journal of Economics*, 105(1), 29–42.
- WILSON, R. B. (1968): “The Theory of Syndicates,” *Econometrica*, 36, 119–132.

TABLE I
Disaster Parameters

	Prior Dist.	Prior Mean	Prior SD	Post. Mean	Post SD.
p_W	Uniform	0.050	0.029	0.037	0.016
p_{CbW}	Uniform	0.500	0.289	0.623	0.076
p_{Cbl}	Uniform	0.050	0.029	0.006	0.003
$1-p_{Ce}$	Uniform	0.500	0.289	0.835	0.027
ρ_z	Uniform	0.450	0.260	0.500	0.034
ϕ	Uniform*	-0.176	0.064	-0.111	0.008
θ	Normal	0.000	0.200	-0.025	0.007
σ_ϕ	Uniform*	0.098	0.047	0.083	0.006
σ_θ	Uniform	0.130	0.069	0.121	0.015

We specify uniform priors on ϕ^* and σ_ϕ^* , the mean and standard deviation of the underlying normal distribution (before truncation). These priors imply (non-uniform) priors on ϕ and σ_ϕ . The numbers in the table refer to the prior mean and standard deviation of ϕ and σ_ϕ .

TABLE II
Disaster Episodes

Country	Start Date	End Date	Max Drop	Perm Drop	Perm/Max	Country	Start Date	End Date	Max Drop	Perm Drop	Perm/Max
Argentina	1890	1908	-0.23	0.02	-0.07	Japan	1940	1952	-0.61	-0.41	0.67
Argentina	1914	1917	-0.13	-0.05	0.37	South Korea	1940	1960	-0.58	-0.48	0.83
Argentina	1930	1933	-0.16	-0.10	0.65	South Korea	1997	2004	-0.23	-0.18	0.81
Argentina	2000	2004	-0.10	-0.01	0.07	Mexico	1914	1918	-0.16	0.27	-1.66
Australia	1914	1923	-0.29	-0.14	0.48	Mexico	1930	1935	-0.24	-0.06	0.23
Australia	1930	1934	-0.24	-0.16	0.65	Netherlands	1914	1919	-0.45	-0.07	0.15
Australia	1939	1956	-0.31	-0.09	0.27	Netherlands	1940	1952	-0.55	-0.10	0.18
Belgium	1913	1920	-0.40	0.05	-0.12	Norway	1914	1924	-0.13	-0.04	0.33
Belgium	1939	1950	-0.52	-0.14	0.26	Norway	1940	1944	-0.08	-0.07	0.84
Brazil	1930	1932	-0.12	-0.05	0.46	Peru	1930	1933	-0.17	-0.08	0.47
Brazil	1940	1942	-0.07	0.00	0.01	Peru	1977	1993	-0.40	-0.37	0.93
Canada	1914	1926	-0.37	-0.20	0.55	Portugal	1914	1921	-0.28	-0.16	0.56
Canada	1930	1933	-0.29	-0.28	0.94	Portugal	1940	1942	-0.09	-0.07	0.74
Chile	1914	1934	-0.53	-0.36	0.69	Spain	1914	1919	-0.10	0.00	0.02
Chile	1955	1958	-0.07	-0.02	0.34	Spain	1930	1961	-0.59	-0.54	0.91
Chile	1972	1987	-0.58	-0.56	0.95	Sweden	1914	1923	-0.21	-0.15	0.72
Denmark	1914	1926	-0.16	-0.08	0.54	Sweden	1940	1951	-0.28	-0.14	0.51
Denmark	1940	1950	-0.28	-0.11	0.40	Switzerland	1914	1921	-0.14	-0.09	0.62
Finland	1890	1893	-0.08	-0.01	0.18	Switzerland	1940	1950	-0.23	-0.15	0.65
Finland	1914	1921	-0.42	-0.22	0.52	Taiwan	1901	1916	-0.24	-0.09	0.37
Finland	1930	1934	-0.23	-0.11	0.49	Taiwan	1940	1955	-0.65	-0.46	0.71
Finland	1940	1945	-0.29	-0.14	0.48	United Kingdom	1914	1921	-0.20	-0.10	0.50
France	1914	1921	-0.22	0.08	-0.36	United Kingdom	1940	1946	-0.20	-0.08	0.39
France	1940	1945	-0.56	-0.07	0.12	United States	1914	1922	-0.24	-0.14	0.57
Germany	1914	1932	-0.45	-0.22	0.48	United States	1930	1935	-0.26	-0.14	0.53
Germany	1940	1950	-0.48	-0.35	0.71						
Italy	1940	1949	-0.33	-0.15	0.45	Average			-0.29	-0.14	0.42
Japan	1914	1918	-0.04	0.12	-2.73	Median			-0.24	-0.10	0.48

A disaster episode is defined as a set of consecudite years for a particular country such that: 1) The probability of a disaster in each of these years is larger than 10%, 2) The sum of the probability of disaster for each year over the whole set of years is larger than 1. Max Drop is the posterior mean of the maximum shortfall in the level of consumption due to the disaster. Perm Drop is the posterior mean of the permanent effect of the disaster on the level potential consumption. Perm/Max is the ratio of Perm Drop to Max Drop.

TABLE III
Mean Growth Rate of Potential Consumption

		Prior		Pre-1946		1946-1972		Post-1973	
	Prior Dist.	Prior Mean	Prior SD	Post. Mean	Post SD.	Post. Mean	Post SD.	Post. Mean	Post SD.
Argentina	Normal	0.02	1.00	0.015	0.010	0.018	0.011	0.008	0.011
Australia	Normal	0.02	1.00	0.014	0.006	0.023	0.005	0.020	0.003
Belgium	Normal	0.02	1.00	0.007	0.006	0.027	0.005	0.019	0.003
Brazil	Normal	0.02	1.00	0.025	0.008	0.037	0.009	0.017	0.008
Canada	Normal	0.02	1.00	0.027	0.005	0.025	0.005	0.018	0.004
Chile	Normal	0.02	1.00	0.019	0.009	0.024	0.009	0.040	0.011
Denmark	Normal	0.02	1.00	0.018	0.004	0.022	0.005	0.012	0.004
Finland	Normal	0.02	1.00	0.025	0.006	0.043	0.007	0.024	0.006
France	Normal	0.02	1.00	0.003	0.003	0.038	0.003	0.019	0.002
Germany	Normal	0.02	1.00	0.014	0.004	0.051	0.005	0.018	0.003
Italy	Normal	0.02	1.00	0.010	0.003	0.046	0.004	0.021	0.003
Japan	Normal	0.02	1.00	0.005	0.004	0.075	0.005	0.022	0.004
Korea	Normal	0.02	1.00	0.017	0.005	0.037	0.010	0.053	0.006
Mexico	Normal	0.02	1.00	0.005	0.008	0.025	0.007	0.016	0.007
Netherlands	Normal	0.02	1.00	0.011	0.004	0.035	0.007	0.016	0.004
Norway	Normal	0.02	1.00	0.015	0.004	0.027	0.004	0.026	0.004
Peru	Normal	0.02	1.00	0.020	0.006	0.030	0.006	0.013	0.008
Portugal	Normal	0.02	1.00	0.017	0.008	0.042	0.007	0.030	0.006
Spain	Normal	0.02	1.00	0.011	0.005	0.055	0.008	0.021	0.004
Sweden	Normal	0.02	1.00	0.026	0.003	0.025	0.004	0.013	0.003
Switzerland	Normal	0.02	1.00	0.013	0.003	0.027	0.003	0.009	0.002
Taiwan	Normal	0.02	1.00	0.007	0.007	0.058	0.009	0.056	0.006
United Kingdom	Normal	0.02	1.00	0.010	0.003	0.020	0.004	0.024	0.003
United States	Normal	0.02	1.00	0.018	0.003	0.025	0.003	0.022	0.003
Median				0.015	0.005	0.029	0.005	0.019	0.004
Simple Average				0.015	0.005	0.035	0.006	0.022	0.005

TABLE IV
Standard Deviation of Non-Disaster Shocks

	Priors			Permanent		Temporary Pre-1946		Temporary Post-1946	
	Dist.	Prior Mean	Prior SD	Post. Mean	Post SD.	Post. Mean	Post SD.	Post. Mean	Post SD.
Argentina	Uniform	0.075	0.04	0.053	0.009	0.020	0.013	0.013	0.009
Australia	Uniform	0.075	0.04	0.017	0.004	0.036	0.008	0.004	0.003
Belgium	Uniform	0.075	0.04	0.020	0.002	0.013	0.009	0.003	0.002
Brazil	Uniform	0.075	0.04	0.047	0.006	0.062	0.011	0.010	0.007
Canada	Uniform	0.075	0.04	0.024	0.003	0.026	0.009	0.003	0.002
Chile	Uniform	0.075	0.04	0.043	0.009	0.038	0.018	0.018	0.010
Denmark	Uniform	0.075	0.04	0.021	0.003	0.005	0.004	0.005	0.003
Finland	Uniform	0.075	0.04	0.031	0.005	0.020	0.008	0.004	0.003
France	Uniform	0.075	0.04	0.014	0.002	0.031	0.005	0.002	0.001
Germany	Uniform	0.075	0.04	0.019	0.002	0.011	0.006	0.002	0.002
Italy	Uniform	0.075	0.04	0.019	0.002	0.011	0.003	0.003	0.002
Japan	Uniform	0.075	0.04	0.022	0.003	0.017	0.005	0.003	0.002
Korea	Uniform	0.075	0.04	0.026	0.004	0.027	0.007	0.004	0.003
Mexico	Uniform	0.075	0.04	0.036	0.004	0.034	0.008	0.005	0.004
Netherlands	Uniform	0.075	0.04	0.023	0.003	0.017	0.006	0.003	0.002
Norway	Uniform	0.075	0.04	0.022	0.002	0.004	0.003	0.004	0.003
Peru	Uniform	0.075	0.04	0.033	0.004	0.007	0.005	0.004	0.003
Portugal	Uniform	0.075	0.04	0.033	0.004	0.023	0.008	0.005	0.003
Spain	Uniform	0.075	0.04	0.024	0.003	0.045	0.008	0.003	0.002
Sweden	Uniform	0.075	0.04	0.019	0.002	0.020	0.004	0.003	0.002
Switzerland	Uniform	0.075	0.04	0.012	0.001	0.039	0.005	0.002	0.001
Taiwan	Uniform	0.075	0.04	0.033	0.004	0.018	0.016	0.004	0.003
United Kingdom	Uniform	0.075	0.04	0.018	0.002	0.003	0.002	0.003	0.002
United States	Uniform	0.075	0.04	0.018	0.002	0.021	0.004	0.003	0.002
Median				0.023	0.003	0.020	0.006	0.003	0.002
Simple Average				0.026	0.004	0.023	0.007	0.005	0.003

TABLE V
Disaster Parameters with Gamma Shocks

	Prior Dist.	Prior Mean	Prior SD	Post. Mean	Post SD.
p_W	Uniform	0.050	0.029	0.035	0.017
p_{CbW}	Uniform	0.500	0.289	0.715	0.094
p_{Cbl}	Uniform	0.050	0.029	0.008	0.004
$1-p_{Ce}$	Uniform	0.500	0.289	0.847	0.029
ρ_z	Uniform	0.450	0.260	0.541	0.037
ϕ	Uniform	0.100	0.058	0.075	0.011
θ	Normal	0.000	0.200	-0.020	0.006
σ_ϕ	Uniform	0.130	0.069	0.091	0.008
σ_θ	Uniform	0.130	0.069	0.110	0.012

TABLE VI
Disasters and the Equity Premium

	Equity Premium	Risk-Free Rate
Baseline	0.048	0.010
No Disasters	0.005	0.042
Permanent, One Period Disasters	0.466	-0.378

All cases have $CRRA = 6.4$, $IES = 2$ and $\beta = \exp(-0.034)$. The return statistics are the log of the average gross return for each asset. The "Equity Premium" is the different between the average return on an unlevered equity claim and bills. The "Risk-Free Rate" is the average return on bills. These results are produced by simulating a long sample from the model with a representative set of disasters.

TABLE VII
Asset Pricing Results for Unleveraged Equity

Specification	CRRA	IES	Full Sample		No Disasters	
			Equity Premium	Risk-Free Rate	Equity Premium	Risk-Free Rate
1. Baseline	6.4	2	0.048	0.010	0.049	0.011
<i>Permanence and Disaster Length:</i>						
2. Permanent	4.4	2	0.048	0.007	0.046	0.015
3. Permanent and One Period	3.0	2	0.048	0.000	0.057	-0.002
<i>Sensitivity to Gamma:</i>						
4. Low Gamma	4.4	2	0.020	0.031	0.020	0.033
5. High Gamma	8.4	2	0.083	-0.017	0.086	-0.019
<i>Model with Gamma Shocks:</i>						
6. Gamma Shocks	6.4	2	0.032	0.022	0.032	0.025
<i>Power Utility:</i>						
7. Power Utility	4.0	0.25	0.012	0.097	-0.011	0.099
8. Power Utility -- One Period/Perm	3.0	0.33	0.048	-0.001	0.060	-0.001
9. Power Utility -- One Period	2.3	0.43	0.048	0.033	0.060	0.009

In all cases, $\beta = \exp(-0.034)$. For case 1, the model of consumption dynamics is parameterized according to the estimates presented in tables 1 through 4. Cases 2-5 and 7-9 are variations on this parameterization. Case 6 is parameterized according to the estimates presented in tables 5 and corresponding estimates of the non-disaster parameters (not-reported). The return statistics are the log of the average gross return for each asset. "Full Sample" refers to a long sample with a representative set of disasters. "No Disaster" refers to a long sample in which agents expect disasters to occur with their normal frequency but non actually occur. The "Equity Premium" is the different between the average return on an unlevered equity claim and bills. The "Risk-Free Rate" is the average return on bills.

TABLE VIII
Long Term Bonds and Parial Default

	(1)	(2)	(3)	(4)
<i>Parameters:</i>				
Coefficient of relative risk aversion	6.4	6.4	7.5	7.5
Intertemporal elasticity of substitution	2	2	2	2
Probability of partial default on perpetuity	0.0	0.4	0.4	0.84
Probability of partial default on one period bond	0.0	0.4	0.4	0.4
<i>Asset Pricing Results:</i>				
Return on one period bond	0.010	0.024	0.014	0.014
Return on perpetuity	-0.021	-0.005	-0.025	0.026
Term premium	-0.032	-0.029	-0.039	0.012
Equity premium	0.048	0.033	0.048	0.048
Average duration of perpetuity in normal times	11.3	9.6	11.9	7.4

In all cases, $\beta = \exp(-0.034)$. The return statistics are the log of the average gross return for each asset.




FIGURE I
A Partially Permanent Disaster

Note: The figure plots the evolution of consumption and potential consumption during and after a disaster lasting six periods with $\rho = 0.6$, $\phi = -0.125$ and $\theta = -0.025$ in each period of the disaster. For simplicity, we abstract from trend growth and assume that all other shocks are equal to zero over this period.




FIGURE II
A Typical Disaster

Note: The figure plots the evolution of log consumption during and after a disaster that strikes in period 1 and lasts for 6 years. Over the course of the disaster, both ϕ and θ take values equal to their posterior means in each period. For simplicity, we abstract from trend growth and assume that all other shocks are equal to zero over this period.






FIGURE III
Ex Ante Disaster Distribution

Note: The solid line is the mean of the distribution of the change in log consumption relative to its previous trend from the perspective of agents that have just learned that they have entered the disaster state but do not yet know the size or length of the disaster. The black dashed line is the median of this distribution. The grey dashed lines are the 5% and 95% quantiles of this distribution.


France



Chile



Korea



United.States




Figure IV
Consumption, Potential Consumption and Disasters in France, Korea, Chile and the U.S.




FIGURE V
World Disaster Probability

Note: The figure plots the posterior mean of $I_{w,t}$, i.e., the probability that the world entered a disaster in each year evaluated using data up to 2006.




FIGURE VI
Asset Prices in Baseline Case with Epstein-Zin-Weil Utility

Note: The figure plots asset returns and detrended log consumption for a “typical” disaster in the baseline case of multi-period disasters with partial recovery when agents have Epstein-Zin-Weil preferences with a coefficient of relative risk aversion of 6.4 and an intertemporal elasticity of substitution of 2. The typical disaster is a disaster that lasts 6 periods and in which the short run and long run disaster shocks take their mean values in each period of the disaster. All other shocks are set to zero.




Figure VII
Posterior Distribution of the Equity Premium




Figure VIII
Distribution of the Coefficient of Relative Risk Aversion




Figure IX
Distribution of the Equity Premium in Data without Disasters




Figure X
Asset Prices in Baseline Case with Power Utility

Note: The figure plots asset returns and detrended log consumption for a “typical” disaster in the baseline case of multi-period disasters with partial recovery when agents have power utility with a coefficient of relative risk aversion of 4. The typical disaster is a disaster that lasts five periods and in which the short run and long run disaster shocks take their mean values in each period of the disaster. All other shocks are set to zero.




FIGURE XI

Asset Prices in Permanent, One Period Case with Power Utility

Note: The figure plots asset returns and detrended log consumption for a “typical” disaster in the case of fully permanent, one-period disasters when agents have power utility with a coefficient of relative risk aversion of 3. The typical disaster is a disaster that lasts one period and in which the short run and long run disaster shocks are equal to -0.40. All other shocks are set to zero.