

Crises and Recoveries in an Empirical Model of Consumption Disasters

Emi Nakamura

Jón Steinsson

Columbia University

Columbia University

Robert Barro

José Ursúa*

Harvard University

Harvard University

January 31, 2010

Abstract

We estimate an empirical model of consumption disasters using a new panel data set on consumption for 24 countries and more than 100 years. 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. We estimate the model using Bayesian methods. Our estimates imply that the probability of entering a disaster is 1.7% per year and that disasters last on average for 6.5 years. In the average disaster episode identified by our model consumption falls by 30% in the short run. In the long run, roughly half of this fall in consumption is reversed. Disasters also greatly increase uncertainty about consumption growth. Our estimates imply a standard deviation of consumption growth during disasters of 12%. We investigate the asset pricing implications of these rare disasters. In a model with power utility and standard values for risk aversion, stocks surge at the onset of a disaster due to agents' strong desire to save. This counterfactual prediction causes a low equity premium, especially in normal times. In contrast, a model with Epstein-Zin-Weil preferences and an intertemporal elasticity of substitution equal to 2 yields a sizeable equity premium in normal times for modest values of risk aversion.

Keywords: Rare disasters, Equity premium puzzle, Bayesian estimation.

JEL Classification: E21, G12

*We would like to thank Timothy Cogley, Xavier Gabaix, Martin Lettau, Frank Schorfheide, Alwyn Young, Tau Zha and seminar participants at various institutions for helpful comments and conversations. Barro would like to thank the National Science Foundation for financial support.

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) argue 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”. One strand of this literature has investigated whether the equity premium may be compensation for the risk of rare but disastrous events. This hypothesis was first put forward by Rietz (1988).¹ A drawback of Rietz’s paper is that it does not provide empirical evidence regarding the plausibility of the parameter values needed to generate a large equity premium based on rare disasters.

Barro (2006) uses data on GDP for 35 countries over the 20th century from Maddison (2003) to evaluate Rietz’s hypothesis empirically. His main conclusion is that a simple model calibrated to the empirical frequency and size distribution of large economic contractions in the Maddison data can match the observed equity premium. In subsequent work, 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. Barro (2006) and Barro and Ursua (2008) analyze the effects of rare disasters on asset prices in a model in which consumption follows a random walk, disasters are modeled as instantaneous, permanent drops in consumption, and the timing of disasters is uncorrelated across countries. They show that it is straightforward to calculate asset prices in this case.

The tractability of the models used in Barro (2006) and Barro and Ursua (2008) comes at the cost of empirical realism in certain respects. First, their model does not allow for recoveries after disasters. Gourio (2008) argues that disasters are often followed by periods of rapid growth. A world in which all disasters are permanent is far riskier than one in which recoveries often follow disasters. Assuming that all disasters are permanent therefore potentially overstates the asset-pricing implications of disasters. Second, their model assumes that the entire drop in consumption due to the disaster occurs over a single time period, as opposed to unfolding over several years as in the data.² Third, their model assumes that output and consumption follow a random walk in normal

¹Other prominent explanations for the equity premium include models with habits (Campbell and Cochrane, 1999), heterogeneous agents (Constantinides and Duffie, 1996) and long run risk (Bansal and Yaron, 2004).

²This assumption is criticized in Constantinides (2008).

times as well as times of disaster. A large literature in macroeconomics has debated the empirical plausibility of this assumption (Cochrane, 1988; Cogley, 1990). Fourth, Barro (2006) and Barro and Ursua (2008) fit their model to the data using an informal estimation procedure based on the average frequency of large economic contractions and the size distribution of peak-to-trough drops in consumption during such contractions. Formal estimation might yield different results.³ Finally, their model does not allow for correlation in the timing of disasters across countries. Relaxing this assumption is important in assessing the statistical uncertainty associated with estimates of the model's key parameters.⁴

In this paper, we consider a richer model of disasters than that considered in Barro (2006) and Barro and Ursua (2008). Our aim is to improve on this earlier work along the dimensions discussed above. Our model allows for permanent and transitory effects of disasters that unfold over multiple years. It allows for transitory shocks to growth in normal times. The model also allows for correlation in the timing of disasters across countries. The model is challenging to estimate using maximum-likelihood methods, because it has a large number of unobserved state variables. It is, however, relatively straightforward to estimate using Bayesian Markov-Chain Monte-Carlo (MCMC) methods. We estimate the model using a Metropolized Gibbs sampler.⁵

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 their characteristics over time and space. We use the Barro-Ursua data on personal consumer expenditure in our analysis.

Our estimates imply that the probability of entering a disaster is 1.7% per year. A 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 then again in the early

³Julliard and Ghosh (2008) develop a novel estimation approach for analyzing the plausibility of rare disasters for the equity premium using data on consumer expenditure and returns for the United States over the twentieth century. They argue that rare disasters are not a likely explanation for the observed equity premium.

⁴If the timing of disasters is assumed to be uncorrelated across countries, the model will interpret the occurrence of, e.g., World War II in a number of countries as many independent observations. This will overstate the statistical precision of the estimates.

⁵A Metropolized Gibbs Sampler is a Gibbs Sampler with a small number of Metropolis steps. See e.g. 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 a comprehensive treatment of these methods.

1980's, and the contraction in South Korea during the Asian financial crisis. On average, disasters last roughly 6.5 years. Consumption drops sharply during disasters. In the disaster episodes we identify in our data, consumption drops on average by 30% in the short run. A large part of this drop in consumption is reversed in the long-run. The long run effect of disaster episodes on consumption in our data is a drop of 14% on average.⁶ Uncertainty about future consumption growth is massive in the disaster state. The standard deviation of consumption growth in this state is roughly 12% per year.

We adopt the representative-agent endowment-economy approach to asset pricing, following Lucas (1978) and Mehra and Prescott (1985). Within this framework, we assess the asset-pricing implications of our model of consumption under alternative specifications of the utility function for the representative consumer. First, consider the case of power-utility analyzed in Mehra and Prescott (1985), Rietz (1988), and Barro (2006). Unlike in the simple disaster model considered in earlier work, the consumption process we estimate generates highly counterfactual implications for the behavior of asset prices during disasters for agents with power utility. For standard parameter values, the onset of a disaster counterfactually generates a stock-market boom, leading to a negative equity premium in normal times.

The key reason for the difference versus earlier models of disasters is that the disasters we estimate unfold over multiple periods rather than occurring instantaneously. Entering the disaster state in our model therefore causes agents to expect steep future declines in consumption. This generates a strong desire to save. When the intertemporal elasticity of substitution is substantially below one, this savings effect dominates the expectations of lower dividends from stocks due to the disaster and therefore raises the price of equity. The large movements in expected consumption growth associated with disasters provide a strong test of consumers' willingness to substitute consumption over time. The strong desire of consumers with low intertemporal elasticities of substitution to smooth consumption over time yields highly counterfactual implications for asset pricing. We interpret the sharp drop in stock prices that typically accompanies the onset of a major disaster as evidence that consumers have a relatively high willingness to substitute consumption over time (at least during disasters).

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.

These results suggest that we should consider a utility specification with a coefficient of relative risk aversion above one and an intertemporal elasticity of substitution above one. We therefore introduce Epstein-Zin-Weil preferences. We find that this model generates asset prices that are much more in line with the data. For an intertemporal elasticity of substitution of 2 and a coefficient of relative risk aversion of 6.5, the model generates an unleveraged equity premium of 4.8%. The asset pricing implications of the model depend importantly on the degree of permanence of the disasters in our estimated model. If the disasters we estimate were completely permanent, the equity premium would be roughly two and a half times larger for the same coefficient of relative risk aversion.

Another way of stating our results is to consider what parameters are appropriate to calibrate a simple model of permanent disasters such as the one considered in Barro and Ursua (2008). As we discuss in section 6, their results can be roughly replicated with $p = 0.0363$ and a constant size of disasters of $b = 0.36$. In contrast, the simple model with constant-sized permanent disasters of size $b = 0.36$ matches our results for a probability of disasters $p = 0.008$. The difference in the two calibrations reflects the fact that we are accounting explicitly for partial recoveries following disasters.

A limitation of some existing models that have been proposed as resolutions of asset-pricing puzzles is that it is difficult to find direct evidence for the underlying mechanism these models rely on. Consequently, the literature has sought to distinguish between these alternative models by assessing whether they resolve not only the equity-premium puzzle but also a number of other asset-pricing puzzles. In contrast, we focus primarily on documenting the existence of rare disasters in long-term consumption data. 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 (2009) embeds disaster risk in a business cycle model.

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

Since rare disasters occur infrequently, by definition, short time series provide little information about the appropriate parameter values for a model of rare disasters. A short sample is likely to contain no disasters even if the true probability is in the range considered by Barro (2006) of 1-2% per year. Furthermore, sample-selection issues are an important problem in disaster studies because data tend to be missing precisely when disasters occur. It is therefore crucial to analyze data covering long time spans, where the starting and ending points are relatively unaffected by the occurrence of disasters.

Barro (2006) uses per capita GDP data from Maddison (2003) for 35 countries for 1900-2000 to estimate the frequency and size distribution of disasters. However, economic models of asset-pricing typically involve consumption, rather than GDP. Barro and Ursua (2008) have since undertaken a major data collection project to develop a long-term panel dataset on per capita personal consumer expenditure—the available data that are likely to proxy well for consumption. They have also substantially modified and extended the Maddison data on GDP. Maddison’s series are in some cases imputed using smooth trends or data from other countries. These imputations tend to occur at times of major upheavals in the countries in question. Barro and Ursua (2008) have removed such imputed observations. In some cases they have been able to locate more comprehensive original sources to fill in gaps.

We use Barro and Ursua’s (2008) dataset on consumer expenditure. Our sample-selection rules follow theirs. We include a country only if uninterrupted data are available back at least before World War I. This procedure yields a sample of 17 OECD countries (4 are dropped because of missing data) and 7 non-OECD countries (11 are dropped due to missing data).⁷ To avoid sample-selection bias problems associated with the starting dates of the series, we include only data after 1890.⁸ The resulting data set is an unbalanced panel for 24 countries, with data from each country starting between 1890 and 1914. The total number of annual observations is 2685.

One limitation of the Barro-Ursua consumption data set is that it does not allow us to distin-

⁷The 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.

⁸Barro and Ursua (2008) use data back to 1870.

guish between expenditures on non-durables and services versus durables. Consumer expenditure generates a flow of consumption services. It is this flow of consumption services that we would like to analyze for asset-pricing purposes. Equating consumption with consumer expenditure may overstate the severity of consumption disasters because consumer expenditure on durables fall more than durable consumption flows. Unfortunately, separate data on durable and non-durable consumption are not available for most of the countries and time periods we study. Barro and Ursua (2008) document the behavior of consumer expenditure on durables and non-durables during consumption disasters for the subset of cases for which data are available. Their analysis shows that declines in consumer expenditure on durables are indeed much larger than declines in consumer expenditure on non-durables during disasters. Nevertheless, declines in consumer expenditure on non-durables are on average only 3 percentage points smaller than for overall consumer expenditure during disasters because durables represent only a small fraction of overall consumer spending. Barro and Ursua (2008) argue that the difference between the decline in overall consumer expenditures versus only non-durable expenditure is even smaller for the case of large disasters. Since durables represent only a small fraction of total expenditure, the contribution of the decline in durable expenditure during a disaster to the overall fall in consumer expenditure is limited (durable expenditure is bounded below by zero).

We also make use of data on the total returns on stocks, bills and bonds. We use the same asset-return data set as Barro and Ursua (2009). These data are based largely on information from Global Financial Data (GFD) but augmented with data from Dimson, Marsh, and Staunton (2002) and other sources. To our knowledge, these are the most comprehensive data sets available on total returns over long periods. Unfortunately, the resulting asset return data are less comprehensive than the Barro-Ursua consumer expenditure and GDP data. For some countries, the returns data start later than the consumer expenditure and GDP data. And for some countries there are gaps in the returns data, often during disaster periods. It may, in principle, be possible to construct a more comprehensive data set from original sources, but such an effort has not yet been undertaken, as far as we know. These data limitations lead us to focus our analysis of asset prices primarily on average returns rather than the evolution of returns over the course of disasters.

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 $\theta_{i,t} \sim N(\theta, \sigma_\theta^2)$. 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, $\phi_{i,t}$ is the short run disaster shock and $\nu_{i,t}$ is an i.i.d. normal shock with a country specific variance $\sigma_{\nu,i}^2$. Since $\theta_{i,t}$ is assumed to affect potential consumption but to leave actual consumption unaffected on impact, it gets added to 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 five periods and in which $\phi_{i,t} = -0.15$ and $\theta_{i,t} = -0.05$ in each period of the disaster. Cumulatively, log consumption drops by roughly 0.60 from peak to trough. Consumption then recovers substantially. In the long run, consumption is 0.25 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 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. 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 1973.⁹ The added flexibility that these assumptions yield dramatically improves the fit of the model to the data. A drawback of modeling these features as one-time events is that whatever generated these features of the data is not part of the process that will generate future consumption in our model. While we have not investigated this issue formally, Bansal and Yaron’s (2004) long-run risk model suggests that persistent movements in the average growth rate of consumption could further raise the equity premium implied by our model.

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 strong 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.

⁹See Perron (1989) and Kilian and Ohanian (2002) for a discussion of trend breaks in macroeconomic aggregates.

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. Maximum likelihood estimation of this model is difficult since it involves carrying out numerical optimization over a high dimensional space of states and parameters. We instead use Bayesian MCMC methods to estimate the model.¹⁰

To carry out our Bayesian estimation we need to specify a set of priors on the parameters of the model. To minimize the influence of the priors on our results, we specify relatively uninformative priors for the majority of the parameters of the model. For a few parameters, however, we specify informative priors. Our main deviation from uninformative priors is that we make assumptions that ensure that “rare disasters” are in fact rare. Specifically, we assume that $p_W \sim U(0, 0.03)$, $p_{CbI} \sim U(0, 0.005)$ and $1 - p_{Ce} \sim U(0, 0.9)$. The first two of these assumptions limit the frequency with which disasters occur to less than 3% per year for world disasters and less than 0.5% per year for idiosyncratic disasters. The third assumption limits the expected length of disasters to less than 10 years.¹¹

The reason we want to focus on estimating rare disasters is that such events can be disproportionately important for asset pricing (Rietz, 1988; Barro, 2006). However, it is important to note that imposing informative priors on p_W , p_{CbI} and p_{Ce} in no way ensures that the rare events we identify will be large disasters. If there are no large disasters in the data, this will be reflected in either a posterior disaster probability close to zero (on which our prior puts substantial weight) or estimates for the means and variances of $\phi_{i,t}$ and $\theta_{i,t}$ implying that the “disasters” are in fact very

¹⁰We sample 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 the baseline 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 alternative specification of our model, in which $\phi_{i,t}$ is assumed to follow a gamma distribution, are based on four independent chains each with 1.2 million draws with the first 150,000 draws from each chain dropped as burn-in. In both cases, 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” (see 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 “overdispersed 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.

¹¹This approach is analogous to the common approach in the asset pricing literature to assume that jumps in returns and volatility are rare and large (Eraker, Johannes and Polson, 2003).

small, which is again a possibility on which our priors put substantial weight. In section 5, we verify that, indeed, if we estimate our model using data generated from a model without disasters, the parameter estimates yield small values for the means and variances of $\phi_{i,t}$ and $\theta_{i,t}$ and a negligible equity premium.

We restrict ρ_z to lie in the interval $[0, 0.9]$. Within this interval, we assume that ρ_z is uniformly distributed. These assumptions ensure that the half-life of the disaster gap is less than 6.5 years. Again, we make these assumptions to ensure that the disasters generated by our algorithm correspond to our intuitive notion of disasters. Our assumptions on ρ_z rule out the possibility that consumption growth in a given period can be explained by disasters that occurred decades earlier.

Finally, we restrict $\nu_{i,t}$ to be small. Specifically, we assume that $\sigma_{\nu,i} \sim U(0, 0.015)$. The only reason we introduce $\nu_{i,t}$ at all to the disaster gap equation is to help the 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.¹²

The priors on all other parameters are very dispersed. In particular, our prior for θ —the mean of the long run disaster shock—is $\theta \sim N(0, 0.2)$. This prior is agnostic about whether disasters have any long run effect at all. Our estimated long run effect thus comes entirely from the data.

Recall that we consider two specifications for the short run shock $\phi_{i,t}$ —a truncated normal distribution and a gamma distribution. For the baseline case of the truncated normal distribution, we assume that $\phi^* \sim U(-0.25, 0)$ and $\sigma_\phi^* \sim U(0.01, 0.25)$. These priors imply a joint prior distribution over ϕ and σ_ϕ . 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 . We assume that $\phi \sim U(-0.25, 0)$ and $\sigma_\phi \sim U(0.01, 0.25)$. These priors imply a joint prior distribution over α_ϕ and β_ϕ .

Our choices for the remaining priors are:

$$\begin{aligned} \sigma_\theta &\sim U(0.01, 0.25), & \mu_{i,t} &\sim N(0.02, 1), \\ \sigma_{\epsilon,i,t} &\sim U(0, 0.15), & \sigma_{\eta,i} &\sim U(0, 0.15), \\ p_{Cbl} &\sim U(0, 1), \end{aligned} \tag{4}$$

¹²MCMC 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.

5 Empirical Results

Table 1 presents our estimates of the disaster parameters for our baseline case, while Tables 2 and 3 present our estimates of $\mu_{i,t}$, $\sigma_{\epsilon,i,t}$ and $\sigma_{\eta,i}$ for this 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.

Consider first our estimates of the disaster parameters in Table 1. We estimate the probability of the start of a world disaster p_W to be 0.021 per year and the probability that a country enters a disaster “on its own” p_{CbI} to be 0.0034 per year. Our estimates for both of these parameters are close to the upper bound of their prior distributions. Our prior restrictions are thus binding for these parameters. The probability that the start of a world disaster will trigger a concurrent disaster in a particular country p_{CbW} is 0.64. The 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. This procedure yields an estimate for the overall probability of 0.017 per year. A centered 90% probability interval for this overall probability is [0.0039, 0.0234]. In contrast, a country that is already in a disaster will continue to be in the disaster in the following year with a 0.847 probability regardless of the world situation. This estimate implies that disasters last on average roughly 6.5 years. Our estimate of ρ_z is 0.495. This value implies that, without further shocks, about half of the disaster gap dissipates each period.

Consumption drops by a large amount on average over the course of the disasters we identify. Our estimate of ϕ —the mean of the short-run shock $\phi_{i,t}$ —is -0.112. In other words, the negative shock to consumption during disasters is on average 11.2% per year. Our estimate of θ —the mean of the long-run shock $\theta_{i,t}$ —is -0.024. This implies that disasters do on average have negative long run effects on consumption. However, the fact that θ is estimated to be much smaller than ϕ implies that a large part of the effect of disasters on consumption in the short run is reversed in the long run.¹³ Our estimate of σ_ϕ and σ_θ —the standard deviation of the short-run shock $\phi_{i,t}$

¹³Davis and Weinstein (2002) document a large degree of recovery at the city level after large shocks.

and long-run shock $\theta_{i,t}$ —are 0.083 and 0.120, 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 the disaster on consumption. The standard deviation of consumption growth in this state is roughly 12% per year.

To get a better sense for what these parameters imply about the nature of consumption disasters, Figures 2 and 3 provide two different visual representations of the size of disasters and the extent of recovery from disasters. Figure 2 plots the impulse response of an “typical disaster”. This prototype lasts for 7 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 five 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 29% fall in consumption (a 0.34 fall in log consumption), while the long-run negative effect of the disaster is approximately 15%.

Figure 3 provides a different view of disasters. Imagine an agent at time 1 who knows that a disaster will begin at time 2 but knows nothing about the character of this disaster beyond the unconditional distribution of disasters. 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 upon receiving news of the disaster relative to what his beliefs were before he received this news. The dashed lines in the figure plot the median and 5% and 95% quantiles of this same distribution. This figure therefore gives an ex ante view of disasters, while Figure 2 gives an ex post view of a particular disaster.

The mean long-run effect of the disaster is similar in magnitude to the long-run effect of the typical disaster depicted in Figure 2. The median long-run effect is smaller than the mean long-run effect because the distribution of disaster sizes is negatively skewed. Figure 2 also shows 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 even 20-25 years later. This long left tail of the disaster distribution is particularly important for asset pricing.

At first glance, Figure 3 seems to tell a different story about the permanence of disasters than the “typical” disaster depicted in Figure 2. The median and mean paths in Figure 3 seem to imply that disasters are much more permanent than is suggested by the typical disaster. This is, however, an artifact of Figure 3 averaging over disasters of different lengths and sizes rather than depicting short-run and long-run effects for any particular disaster. Disasters of different lengths will reach

their troughs at different points in time—for 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 towards 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.

Table 4 reports summary statistics for the main disaster episodes identified by our model. Our Bayesian estimation procedure does not deliver a definitive judgement on whether a disaster occurred at certain times and places but rather provides a posterior probability of whether a disaster occurred. 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.¹⁴ Using this definition, we identify 50 disaster episodes. These episodes vary greatly in size and shape. On average, the maximum drop in consumption due to the disasters is 30%. The permanent effect of disasters on consumption is on average 14%. However, the largest short run effects of a disaster are 66% and 61% drops in consumption in Taiwan and Japan, respectively, during World War II.¹⁵ Quite a few of the disaster episodes have huge long-run effects. For example, we estimate the long-run effect of the disaster episode in Chile in the 1970's and early 1980's to be a 56% drop in consumption relative to a counterfactual consumption path with no disaster, while the long-run effect of the Spanish Civil War and the subsequent turmoil was a 53% drop in consumption relative to a counterfactual no-disaster consumption path.

The bulk of the disaster episodes we identify occur during World War I, the Great Depression

¹⁴More 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.

¹⁵In all cases, these statistics measure the negative effect of the disaster on the level of consumption relative to the counterfactual scenario where the country instead experienced normal trend growth.

and World War II. Figure 4 plots our estimates of the probability that a “world disaster” began in each year.¹⁶ Our model clearly identifies 1914, 1930 and 1940 as years in which world disasters began. The model also identifies five disaster episodes that are not associated with world disasters. These include one of the most serious disaster episodes—Chile during the early period of the reign of Augusto Pinochet. In a few cases, our model is not able 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. An example of this is the Great Depression and the Spanish Civil War in Spain.

Figure 5 provides more detail about how our model interprets the evolution of consumption for France, Korea, Chile and the United States.¹⁷ 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. The left axis gives values of the probability of disaster, while the right axis gives values for log consumption and potential consumption.

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 1957 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 (46%). This pattern is typical of the experience of Asian countries in our sample during WWII. For South Korea, we also identify the Asian Financial Crisis as a disaster.¹⁸

Chile is one of the most volatile countries in our sample. Our model identifies three disaster episodes for Chile. The first begins in WWI and spans the early years of the Great Depression. The second disaster occurred in 1955-59 during the final reign of Carlos Ibanez. The third disaster in Chile began in 1970 during the tenure of Salvador Allende but intensified greatly in the early years of Augusto Pinochet’s rule. The late 1970’s and early 1980’s are a period of recovery. But

¹⁶This 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?

¹⁷More detailed figures for all the countries in our study are reported in a web appendix.

¹⁸Countries such as Indonesia and Thailand that likely also experienced disasters during the Asian Financial Crisis, are not in the data set.

another period of huge decline in consumption starts in 1982 at the time of the Latin American debt crisis and lasts until 1987. This long disaster period—from 1973 to 1987—is the most severe disaster episode we identify outside of periods of major world wars.

The last panel in Figure 5 plots results for the United States. Relative to most other countries in our sample, the United States was a tranquil place during our sample period. The model identifies two disaster episodes for the U.S. 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 the 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.

According to our model, there have been no world disasters since the end of WWII.¹⁹ It is natural to ask whether this pattern provides evidence against the model. In fact, the rare nature of world disasters implies that the posterior probability of experiencing no world disasters over a 61 year stretch is roughly 27%. One could also ask whether the relative tranquility of the United States experience since the Great Depression provides evidence that the U.S. is fundamentally different from other countries in our sample. However, the posterior probability for a randomly selected country experiencing no disasters over a 73-year stretch is 0.31 according to our model. The posterior probability of at least one out of 24 countries experiencing no disaster over a 73-year stretch is 0.88. Therefore, the tranquility of the U.S. experience (which is not randomly selected) does not provide evidence against our model.

Tables 2 and 3 present the remaining parameter estimates for our empirical model. Table 2 presents country-specific estimates of the mean growth rate of potential consumption for the countries in our sample. Recall that our model allows for breaks in the mean growth rate of potential consumption in 1946 and 1973. We estimate sizable breaks of this kind in 1946 and 1973 for many countries, especially those for which WWII was most disastrous.²⁰

¹⁹Our sample period ends in 2006 and thus does not cover the current world-wide recession, which may in the future be identified as a world disaster.

²⁰If we do not allow for breaks in the mean growth rate, the model yields quite different results and a substantially worse overall fit to the data. In this case, WWII is interpreted as having had a very large positive long-run effect for Japan, France and Italy (but a negative short-run effect). The model interprets the rapid growth from 1946 to 1973 as slow convergence to a much higher level of potential consumption. The notion that disasters will generally lead to large long-run increases in the level of consumption does not line up well with other disasters in our data—e.g. WWI. An alternative model that we find more plausible is that changes in the mean growth rate ($\mu_{i,t}$ in our model) tend to be particularly large during disasters. Such persistent changes in the mean growth rate of consumption are the focus of the long-run risk model of Bansal and Yaron (2004). Incorporating this feature is beyond the scope of the current paper. Instead, we allow for break in $\mu_{i,t}$ in 1946 and 1973. Including these breaks substantially improves the fit of the model relative to the specification without breaks. Modeling long-run risk could raise the equity premium

Table 3 presents country-specific estimates of the variances of the permanent and transitory shocks to consumption. We allow for a break in the variance of the transitory shock in 1946. We find a great deal of evidence for such a break. For all but four 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.

As we discuss earlier in the paper, we also consider a model with gamma-distributed short-run shocks. Like the truncated normal distribution we use in our baseline model, the gamma distribution is one-sided. However, the gamma distribution has excess kurtosis relative to the truncated normal distribution. Table 5 presents our estimates of the main disaster parameters when we assume that $\phi_{i,t}$ has a Gamma distribution. Most of the estimates are quite similar to the baseline case. The overall probability that a country enters a disaster in a given year is 1.8% versus 1.7% in the baseline model. The mean of the long run shock is very similar at -2.0% and so is the persistence of the disaster gap at 0.56. The main difference that arises is that the gamma model assigns a somewhat larger portion of the volatility of consumption during disasters to the short-run shock as oppose to the long-run shock. The standard deviations of these shocks are 0.095 and 0.108, respectively, while they are 0.083 and 0.120, respectively, in the baseline model.

6 Asset Pricing

We follow Mehra and Prescott (1985) in analyzing the asset pricing implications of the consumption process we estimate above 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:

$$E_t \left[\beta^\xi \left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{(-\xi/\psi)} R_{w,t,t+1}^{-(1-\xi)} R_{i,t,t+1} \right] = 1, \quad (5)$$

where $R_{i,t,t+1}$ denotes the gross return on an arbitrary asset in country i from period t to period $t + 1$, $R_{w,t,t+1}$ denotes the gross return on the agent's wealth, which in our model equals the endowment stream. The parameter β represents the subjective discount factor of the representative implied by our model by increasing the overall risk faced by investors.

consumer. The parameter $\xi = \frac{1-\gamma}{1-1/\psi}$, where γ is the coefficient of relative risk aversion and ψ is the intertemporal elasticity of substitution (IES).²¹

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 is equal to the reciprocal of the IES— $\gamma = 1/\psi$. In other words, a single parameter governs consumers’ willingness to bear risk and willingness to substitute consumption over time. Bansal and Yaron (2004) and Barro (2009), among others, have emphasized the importance of delinking these two features of consumer preferences. Our results below provide additional evidence in support of a more flexible model.

The asset-pricing implications of our model with Epstein-Zin-Weil (EZW) preferences cannot be derived analytically.²² We therefore use standard numerical methods.²³ Initially, we calculate returns for two assets: a one period risk-free bill and an unleveraged claim on the consumption 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.

Differences in the discount factor have only minimal effects on the equity premium in our model.²⁴ These differences do, however, affect the risk-free rate. Given a calibration of γ and ψ , we can pick β to match the risk-free rate generated by the model to the risk-free rate observed in the data. We set $\beta = \exp(-0.034)$ to match the risk-free rate for our baseline specification. The equity premium is highly sensitive to the coefficient of relative risk aversion γ in our model. We present results for a range of values for γ that includes the value that matches the average equity premium in the data.

²¹The 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. See also Rubinstein (1974) and Caselli and Ventura (2000). Constantinides and Duffie (1996) argue that highly persistent, heteroscedastic uninsurable income shocks can resolve the equity-premium puzzle.

²²An analytical solution for asset prices may be derived in the case of permanent disasters (Barro, 2009).

²³We solve the integral in equation (5) on a grid. Specifically, we start by solving for the price-dividend ratio for a consumption claim. In this case we can rewrite equation (5) 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 (5) for other assets as $PDR_t = E_t[f(\Delta C_{t+1}, \Delta D_{t+1}, PDR_t^C, PDR_t)]$, 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 approach used by Campbell and Cochrane (1999) and Wachter (2008).

²⁴In the continuous time limit of our discrete time model, the equity premium is unaffected by β .

There is a debate in the macroeconomics and finance literature about the appropriate parameter value for the IES. Hall (1988) estimates the IES to be close to zero. His estimates of the IES are obtained by analyzing the response of aggregate consumption growth to movements in the interest rate over time. Yet, as noted by Bansal and Yaron (2004) and Gruber (2006), such estimates are potentially subject to important endogeneity concerns. The interest rate and consumption growth are results 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 of this dispersion in empirical estimates, 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 (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 $IES = 2$ as our baseline case.

Barro and Ursua (2008) present data on rates of return for stocks, bonds and bills for 17 countries over long periods. The average arithmetic real rate of return on stocks in their data is 8.1% per year. The average arithmetic real rate of return on short term bills is 0.9% per year. The average equity premium in their data is therefore 7.2% per year. If we view stock returns as a leveraged 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 ($7.2/1.5$).²⁵ We therefore take 4.8% per year as the target

²⁵Dividing the equity premium for leveraged equity by one plus the debt-equity ratio to get a target for unleveraged equity is exactly appropriate in the simple disaster model of Barro (2006). Abel (1999) argues for approximating leveraged equity by a scaled consumption claim. Bansal and Yaron (2004) and others have adopted this approach. For our model, these two approaches yield virtually indistinguishable results.

for our analysis.

In our analysis, we use the actual consumption data for the countries we study. These data presumably 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. 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.²⁶

6.1 The Equity Premium with Epstein-Zin-Weil Preferences

The asset-pricing results for our model are presented in Table 6. Specification 1 in the table is for our baseline model and preferred preference parameters. The table also presents results for a number of alternative specification. 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 in most OECD countries. The statistics we report are the logarithm of the arithmetic average gross return on each asset ($\log E[R_{i,t,t+1}]$).

For the baseline model with $IES = 2$, we present results for three levels of risk aversion: $\gamma = 4.5, 6.5, 8.5$. With $\gamma = 6.5$, the model generates an unleveraged equity premium of 4.9% per year. Our baseline model with an $IES = 2$ and $\gamma = 6.5$ thus matches the equity premium and the risk-free rate observed in the data. The model generates an equity premium of 1.8% when $\gamma = 4.5$, while it generates an equity premium of 8.7% when $\gamma = 8.5$. This illustrates the high degree of sensitivity of the equity premium to the degree of risk-aversion in the presence of disasters. For comparison, specification 6 in Table 6 presents results for a version of the model in which we turn off the disasters completely. Without disasters, the model generates an equity premium that is too small by a factor of 10. This finding is in line with Mehra and Prescott (1985).

²⁶Papers 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).

Figure 6 depicts equity and bond returns over the course of a “typical” disaster when the $IES = 2$ and $\gamma = 6.5$. When the news arrives that a disaster has struck, the stock market crashes. In contrast, bills are risk-free in the short run. Their returns are thus not affected in this initial period. The crash in the stock market at the onset of the disaster coincides with a large drop in consumption. Stocks must yield a considerable return-premium over bills in normal times to compensate for the risk of a disaster.

When the news arrives that the disaster has ended, stocks surge relative to bills. With an $IES = 2$, the crash at the beginning of the disaster and the boom at the end of the disaster are roughly equally large and the average equity premium during the disaster (excluding the initial crash and the boom at the end) is similar to the average equity premium in normal times. The average equity premium in a long sample with a representative number of disasters is thus similar to the average equity premium in a long sample for which agents expect disasters to occur with their normal frequency but no disasters actually occur (the case considered in the three rightmost columns of Table 6).

An important innovation of our model relative to earlier models of disasters is that we allow for partial recovery after disasters. To assess the importance of this feature of our model, specification 2 in Table 6 presents results for a case in which disasters are completely permanent on average as in Barro (2006) and Barro and Ursua (2008). This case generates an equity premium of 12.6%. Specification 3 in Table 6 presents results for the polar opposite case of completely temporary disasters. In this case, the equity premium falls to 0.6%. A world in which disasters are completely permanent is clearly much riskier than a world in which there is substantial recovery after disasters. The partial recoveries that we estimate after disasters is the main reason why we match the equity premium with a value of 6.5 for the coefficient of relative risk aversion, while Barro and Ursua (2008) can match the equity premium with a value of 3.5.²⁷

Specification 8 in Table 6 is for the case in which the short run disaster shocks follow a (one sided) Gamma distribution. With $\gamma = 6.5$ and an IES of 2, the equity premium is 3.4% and the risk-free rate is 2.0%. The gamma model matches the equity premium and risk-free rate when $\gamma = 7.5$. 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 than the baseline model.

²⁷In the specification in which disasters are completely permanent on average we assume that $\theta_{i,t} \sim N(\phi, \sigma_\phi)$. In the specification in which disasters are completely temporary we set $\theta_{i,t} = 0$.

The asset-pricing exercises discussed above are based on the posterior means of the parameters of our model. These calculations thus ignore sampling error in our parameter estimates. Given the limited amount of data we have to estimate the frequency, size, and shape of rare disasters, the posterior standard deviation of the parameters governing disasters are in some cases substantial. Using the posterior distribution of the parameters of our model, we can calculate a posterior distribution for the equity premium. This allows us to investigate the precision of our model's implications for the equity premium. The posterior distribution for the equity premium implied by the posterior distribution of the parameters of our model is plotted in Figure 7. In calculating this distribution, we assume that agents have $\beta = \exp(-0.034)$, $\gamma = 6.5$ and $\psi = 2$. Figure 7 shows that our estimates place a large weight on parameter combinations that generate a significant equity premium. Parameter combinations that generate an equity premium of less than 3% get less than 3% weight. The centered 90% probability interval for the equity premium is $[0.032, 0.073]$.

An important question is whether our results regarding the equity premium are somehow “built in” to our prior or estimation algorithm. To assess the degree to which our results on the equity premium are driven by the data as opposed to our priors and estimate algorithm, we 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 8. For this alternative data set, our model places a large probability (roughly 77%) on the equity premium being below 1%. Clearly, our results would be very different if there in fact were no disasters in the data. The distribution has a long right tail reflecting the fact that even a data set the size of ours with no disasters would not entirely rule out the possibility that such events could occur.

During the disaster, consumers expect consumption to keep falling and thus have an incentive to save. This drives up the price-dividend ratio of assets and drives down their expected returns. As a consequence, stock and bill returns are lower on average during disasters than during normal times even after the initial crash (see Figure 6). Furthermore, the return on stocks and bills is temporarily high during the recovery period after a disaster. 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 have been 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 after WWI.

In our model, consumption growth is predictable following major disasters, as the country recovers to the level of potential consumption. Outside of disasters and their immediate aftermath, however, consumption growth is hard to forecast. The explanatory power of the price dividend ratio in predicting future consumption growth at both medium and short horizons is close to zero.²⁸ In this regard, our model differs from the long-run risks model of Bansal and Yaron (2004) which 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 the U.S. consumption data, particularly in the post-WWII period.

One way to think about the importance of the features we have added relative to earlier work on disasters is to ask how we could recalibrate the simpler model used in Barro and Ursua (2008) to generate an equity premium of the same size as the one our model yields. Recall that, in their model, all disasters are completely permanent. The appropriate parameters for the model may, therefore, differ from the observed probability and size of (partially transient) disasters in the data. The equity premium in Barro and Ursua (2008) is given by

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

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. For simplicity, consider a version of this model in which b is a constant. Barro and Ursua's (2008) results can be roughly replicated with $p = 0.0363$ and a constant size of disasters of $b = 0.36$. In contrast, the simple model with constant-sized permanent disasters of size $b = 0.36$ matches our results for a probability of disasters $p = 0.008$. The substantially lower probability of disasters in our calibration—and the corresponding lower equity premium for any given value of risk aversion—arises mainly from the fact that our estimated model includes partial recoveries after many disasters.

²⁸Specifically, we have analyzed regressions of consumption growth at one, 3 and 5 year horizons on the current price dividend ratios. The R^2 of such regressions is consistently 3% or less.

6.2 The Equity Premium with Power Utility

As we discussed before, there is no consensus in the literature regarding the appropriate value for the intertemporal elasticity of substitution. In our baseline results, we follow Bansal and Yaron (2004) and Barro and Ursua (2008) in assuming that the IES is larger than one. To see why this is important in our context, it is instructive to consider asset-pricing results for our model when the representative consumer has power utility. Results for the case of $\gamma = 1/\psi = 4$ are presented in specifications 4 and 5 of Table 6—this is the utility specification used by Barro (2006).

For the baseline model with partially temporary multi-period disasters, the equity premium is 0.8%. This is only about 60% of the equity premium the model generates with $\gamma = 4$ and $\text{IES} = 2$. However, a more serious concern is that, conditional on no disasters, the equity premium is -0.9%, i.e., lower than in a model in which no disasters can happen. The overall equity premium is, therefore, coming entirely from superior equity returns *during* disasters and the equity premium in normal times is negative. 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 important difference is that disasters unfold over multiple periods in our model. Specification 5 of Table 6 presents results for a version of our model with single-period, fully permanent disasters and power utility parameterized to match the equity premium of 3.6% in Barro (2006).²⁹ The equity premium conditional on no disasters in this specification is 4.5%.

To provide intuition for these results, Figure 9 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 for this large positive return is that entering the disaster state causes agents in the model to expect further drops in consumption going forward. Since the agents in the model have an IES equal to only 1/4 they have a tremendous desire to smooth consumption over time and, hence, have a tremendous desire to save when they receive news that they are in a disaster. This desire to save is so strong that it dominates the fact that entering a disaster is bad news about the dividends on stocks. The disaster therefore

²⁹We make σ_{ϕ} and σ_{θ} small and set $\phi = \theta = 0.44$.

causes a sharp rise in stock prices.³⁰ In contrast to stocks, the one period risk-free bond delivers a “normal” return in the first period of the disaster. Together, these two facts imply that agents do not demand a high return for holding stocks in normal times as a compensation for disaster risk. However, stockholders demand a large equity premium during disasters partly as compensation for the “risk” that the disaster might end, lowering the demand for assets and causing a sharp drop in stock prices.

Figure 10 presents a set of analogous results for the case of a single-period permanent disaster with power utility. The results for this case are much more intuitive. In this case, the disaster occurs instantaneously with no change in expected consumption growth going forward. 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. Needless to say, the prediction of our multi-period disaster model with power utility—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.

In the one-period disaster—as in the Barro-Rietz model—the entire drop in consumption coincides with the crash of the stock market. In contrast, in the multi-period disaster, the largest movement in the stock market occurs at the onset of the disaster—because that is when much of the news associated with the disaster is revealed—but the fall in consumption occurs over several years. The equity premium is larger the larger is the coincidence of timing of the drop in consumption and the crash in the stock market. The high degree of coincidence of the drop in consumption and the crash in the stock market in the one-period disaster is thus one factor that contributes to the large equity premium in this case.

6.3 Long Term Bonds, Inflation Risk and Partial Default

Our model generates variations in the expected growth rate of consumption. This property leads to a non-trivial term structure of real interest rates. While one period bonds are risk-free over a one-period horizon, the same is not true of longer-term bonds. The return on long-term bonds relative to short-term bills will reflect their relative risk properties. Most long-term government

³⁰Similar 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.

bonds promise a fixed set of payments in nominal terms. Inflation risk will therefore affect the prices of these securities. Barro and Ursua (2008) present data on real returns on long-term bonds for 15 countries over long sample periods. The underlying claims are nominal government bonds usually of around ten-year maturity. The average arithmetic real rate of return on these bonds in their data 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 their data is 1.2% per year. To approximate such 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 ten year coupon bonds.³¹ Importantly, the coupon payments are assumed to be set in real terms and are risk free.

The returns on such long-term bonds in the baseline model are reported in Table 7. The average return is -2.3% per year. This 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 bond has such a low average return in the presence of disasters is that it is 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. This response 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 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.

Barro and Ursua (2009) consider peak to trough drops in stock prices over time periods that correspond roughly to consumption disasters. They note that in many cases large negative stock returns occur slightly prior to or coincident with the consumption disasters. In some cases, however, the timing of negative stock returns differs relative to the timing of the consumption disaster. They

³¹The 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%.

argue that measurement issues, including effect from price controls, can help to explain this variation in timing. Therefore, they advocate a flexible timing methodology for computing covariances.³²

Extending the asset price calculations of Barro and Ursua (2009) to bills, we find that in 74% of the largest consumption disasters—25 cases out of 34—stock returns are lower than bill returns.³³ 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 returns is most suspect because of market closure and controls on goods and asset prices.

We follow Barro (2006) in calculating asset prices under the assumption that bills and bonds experience partial default with a specified probability during disasters due to large amounts of inflation. Specifically, we assume that with some probability the return on bills and bonds is equal to the return on equity. The calculations discussed above suggest that an appropriate calibration of the probability of partial default is 26%. An earlier calculation by Barro (2006) based on somewhat less extensive data suggested a probability of 40% for partial default. To be conservative, we also consider a case in which the probability of partial default is 40%.

The first column of Table 7 restates our baseline result from Table 6. The second and third columns of Table 7 report results for calibrations that allow for partial default on bills. For $\gamma = 6.5$, this modification lowers the equity premium from 4.8% to 3.4%. Raising the coefficient of relative risk aversion 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 an immediate jump in the price level. The fourth column in Table 7 reports results for a case in which the probability of partial default on the perpetuity is 40% but there is no partial default on bills. In this case, the average return on the perpetuity is -0.7% implying a term premium of -1.6%. If long term bonds experience

³²The appropriate timing assumption when calculating the covariance between returns and changes in consumption for asset pricing purposes depends on the nature of preferences. In a model with habit formation, for example, the timing of an increase in marginal utility does not coincide directly with the timing of movements in consumption. This type of model is beyond the scope of our current analysis.

³³We consider the subset of events in which the peak-to-trough drop in consumption is larger than 17%. This procedure yields a set of events that corresponds closely to the disaster episodes in our model. For the subset of countries that we use to estimate our model, we get 48 events as compared to 50 disaster episodes identified by our model. The average drop in consumption for these events is 32%, compared to 30% for our disaster episodes. There are 34 events for which we have data on both stock and bill returns.

partial default in an even larger fraction of disasters, the term premium will be higher. We can thus match the term premium by raising the probability of partial default. The fifth column presents results for a case which matches the term premium of 1.2% with a probability of partial default of 73.5%. In other words, to match the term premium in the data, the perpetuity we consider need only provide insurance against roughly one of every four disasters.

Our model generates the implication that, without default risk on bonds (for example associated with inflation risk for nominal bonds), the term structure is downward-sloping. Introducing risk of partial default on longer term bonds allows us to match the fact that the nominal term structure is upward-sloping in the Barro-Ursua data. If most of the default risk comes from inflation risk, our model implies that the term structure on 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. They caution that this evidence is hard to assess because of the short sample and poor liquidity in the TIPS market.³⁴

7 Conclusion

In this paper, we estimate an empirical model of macroeconomic disasters, building on the work of Rietz (1988), Barro (2006) and Barro and Ursua (2008). The key innovations of our model are that we allow disasters to be partly transitory, to unfold over multiple periods, and for the timing of disasters to be correlated across countries. Furthermore, we use a formal Bayesian estimation procedure to match the data to the model. We find that it is possible to match the observed equity premium using the estimated representative-agent model, with a coefficient of relative risk aversion of 6.5 and an intertemporal elasticity of substitution of 2. The degree of risk aversion needed to match the empirical equity premium in our model is higher than in Barro (2006) and Barro and Ursua (2008) mainly because we estimate substantial recoveries after disasters.

Our asset-pricing results depend on the assumption that, unlike under power utility, the co-

³⁴See also Evans (1998), Barr and Campbell (1997) and Campbell, Shiller, and Viceira (2009).

efficient of relative risk aversion need not equal the reciprocal of the intertemporal elasticity of substitution (IES). Specifically, with Epstein-Zin-Weil preferences, the risk-aversion coefficient and the IES can both exceed one. With an intertemporal elasticity of substitution substantially below one, our asset-pricing model counterfactually generates stock market booms at the onset of disasters.

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 $f(\Phi_i|\Phi_{-i}, Y)$.³⁵

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 $c_{i,t}^c = c_{i,t} - c_{i,t}^t$. We set $z_{i,t} = \min(\max(-0.5, c_{i,t}^c), 0)$ 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 $z_{i,t}$.

³⁵The 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} .

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 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.
- 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 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. KLESCHCHLSKI, 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.

- CASELLI, F., AND J. VENTURA (2000): “A Representative Consumer Theory of Distribution,” *American Economic Review*, 90, 909–926.
- 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.
- GOURIO, F. (2008): “Disasters and Recoveries,” *American Economic Review*, 98(2), 68–73.
- (2009): “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.
- JULLIARD, C., AND A. GHOSH (2008): “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.
- 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.
- PERRON, P. (1989): “The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis,” *Econometrica*, 57, 1361–1401.
- 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.
- 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.
ρ_W	Uniform	0.015	0.009	0.021	0.006
ρ_{CbW}	Uniform	0.500	0.289	0.643	0.074
ρ_{Cbl}	Uniform	0.0025	0.0014	0.0034	0.0010
$1-p_{Ce}$	Uniform	0.450	0.260	0.847	0.023
ρ_z	Uniform	0.450	0.260	0.495	0.040
ϕ	Uniform*	-0.176	0.064	-0.112	0.008
θ	Normal	0.000	0.200	-0.024	0.007
σ_ϕ	Uniform*	0.098	0.047	0.083	0.006
σ_θ	Uniform	0.130	0.069	0.120	0.016

We specify uniform priors on ϕ^* and σ_ϕ^* . 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
Mean Growth Rate of Potential Consumption

	Prior Dist.	Prior		Pre-1946		1946-1972		Post-1973	
		Prior Mean	Prior SD	Post. Mean	Post SD.	Post. Mean	Post SD.	Post. Mean	Post SD.
Argentina	Normal	0.02	1.00	0.016	0.010	0.017	0.011	0.007	0.010
Australia	Normal	0.02	1.00	0.014	0.006	0.022	0.005	0.020	0.003
Belgium	Normal	0.02	1.00	0.007	0.006	0.027	0.005	0.019	0.004
Brazil	Normal	0.02	1.00	0.024	0.008	0.037	0.009	0.017	0.008
Canada	Normal	0.02	1.00	0.027	0.005	0.026	0.005	0.018	0.004
Chile	Normal	0.02	1.00	0.018	0.009	0.024	0.010	0.040	0.011
Denmark	Normal	0.02	1.00	0.018	0.004	0.021	0.005	0.012	0.004
Finland	Normal	0.02	1.00	0.025	0.006	0.042	0.007	0.023	0.006
France	Normal	0.02	1.00	0.004	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	0.006	0.035	0.009	0.052	0.006
Mexico	Normal	0.02	1.00	0.005	0.007	0.025	0.007	0.015	0.006
Netherlands	Normal	0.02	1.00	0.011	0.004	0.034	0.007	0.015	0.004
Norway	Normal	0.02	1.00	0.015	0.004	0.028	0.004	0.025	0.004
Peru	Normal	0.02	1.00	0.020	0.006	0.030	0.006	0.013	0.009
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	0.028	0.003	0.009	0.002
Taiwan	Normal	0.02	1.00	0.007	0.007	0.057	0.009	0.055	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	0.004	0.022	0.003
Median				0.014	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 III
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	0.054	0.007	0.022	0.016	0.012	0.009
Australia	Uniform	0.075	0.04	0.018	0.004	0.036	0.008	0.003	0.002
Belgium	Uniform	0.075	0.04	0.020	0.003	0.013	0.009	0.003	0.002
Brazil	Uniform	0.075	0.04	0.047	0.006	0.062	0.011	0.011	0.007
Canada	Uniform	0.075	0.04	0.024	0.003	0.027	0.008	0.003	0.002
Chile	Uniform	0.075	0.04	0.043	0.009	0.036	0.018	0.018	0.011
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	0.004	0.019	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	0.010	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	0.027	0.004	0.027	0.007	0.004	0.003
Mexico	Uniform	0.075	0.04	0.037	0.004	0.034	0.008	0.005	0.004
Netherlands	Uniform	0.075	0.04	0.024	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	0.024	0.008	0.005	0.004
Spain	Uniform	0.075	0.04	0.024	0.004	0.046	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	0.001	0.001
Taiwan	Uniform	0.075	0.04	0.033	0.003	0.015	0.014	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	0.021	0.006	0.003	0.002
Simple Average				0.026	0.004	0.023	0.007	0.005	0.003

TABLE IV
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	1907	-0.22	0.02	-0.10	South Korea	1997	2005	-0.23	-0.18	0.77
Argentina	1914	1917	-0.12	-0.05	0.37	Mexico	1914	1918	-0.16	0.27	-1.76
Argentina	1930	1933	-0.15	-0.10	0.63	Mexico	1930	1935	-0.24	-0.05	0.20
Australia	1914	1922	-0.29	-0.14	0.50	Netherlands	1914	1919	-0.45	-0.07	0.16
Australia	1930	1934	-0.24	-0.16	0.64	Netherlands	1940	1951	-0.55	-0.08	0.14
Australia	1939	1955	-0.31	-0.07	0.22	Norway	1914	1924	-0.13	-0.04	0.28
Belgium	1913	1920	-0.39	0.06	-0.15	Norway	1940	1944	-0.07	-0.06	0.87
Belgium	1940	1950	-0.51	-0.12	0.24	Peru	1930	1933	-0.16	-0.08	0.47
Brazil	1930	1933	-0.11	-0.06	0.50	Peru	1977	1993	-0.39	-0.36	0.92
Canada	1914	1926	-0.37	-0.19	0.52	Portugal	1914	1921	-0.28	-0.15	0.55
Canada	1930	1933	-0.29	-0.27	0.93	Portugal	1940	1942	-0.09	-0.07	0.74
Chile	1914	1934	-0.53	-0.36	0.68	Spain	1914	1919	-0.10	0.00	0.05
Chile	1955	1959	-0.06	-0.02	0.37	Spain	1930	1961	-0.58	-0.53	0.91
Chile	1970	1987	-0.58	-0.56	0.95	Sweden	1914	1923	-0.21	-0.15	0.71
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	0.38	Switzerland	1914	1921	-0.14	-0.09	0.63
Finland	1914	1920	-0.42	-0.22	0.53	Switzerland	1940	1950	-0.23	-0.15	0.68
Finland	1930	1935	-0.23	-0.11	0.49	Taiwan	1901	1915	-0.24	-0.08	0.33
Finland	1940	1945	-0.29	-0.14	0.49	Taiwan	1940	1955	-0.66	-0.45	0.69
France	1914	1921	-0.22	0.08	-0.35	United Kingdom	1914	1921	-0.21	-0.11	0.51
France	1940	1945	-0.56	-0.07	0.13	United Kingdom	1940	1946	-0.20	-0.07	0.35
Germany	1914	1932	-0.45	-0.22	0.48	United States	1914	1922	-0.25	-0.14	0.58
Germany	1940	1949	-0.48	-0.34	0.71	United States	1930	1935	-0.26	-0.14	0.53
Italy	1940	1949	-0.33	-0.15	0.46						
Japan	1914	1918	-0.04	0.12	-2.76	Average			-0.30	-0.14	0.37
Japan	1940	1951	-0.61	-0.41	0.68	Median			-0.25	-0.11	0.50
South Korea	1940	1957	-0.57	-0.46	0.80						

A disaster episode is defined 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%, 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 V
Disaster Parameters with Gamma Shocks

	Prior Dist.	Prior Mean	Prior SD	Post. Mean	Post SD.
ρ_W	Uniform	0.015	0.009	0.020	0.006
ρ_{CbW}	Uniform	0.500	0.289	0.724	0.093
ρ_{Cbl}	Uniform	0.0025	0.0014	0.0035	0.0010
$1-\rho_{Ce}$	Uniform	0.450	0.260	0.862	0.021
ρ_z	Uniform	0.450	0.260	0.555	0.039
ϕ	Uniform	-0.100	0.058	0.084	0.009
θ	Normal	0.000	0.200	-0.020	0.006
σ_ϕ	Uniform	0.130	0.069	0.095	0.008
σ_θ	Uniform	0.130	0.069	0.108	0.013

TABLE VI
Asset Pricing Results

Specification	Coef. Rel. Risk Av.	Int. Elast. Sub.	Full Sample			Cond. On No Disasters		
			Equity Return	Bill Return	Equity Premium	Equity Return	Bill Return	Equity Premium
1. EZW Utility -- Baseline	4.5	2	0.050	0.032	0.018	0.051	0.034	0.017
	6.5	2	0.058	0.009	0.049	0.058	0.010	0.048
	8.5	2	0.065	-0.022	0.087	0.066	-0.024	0.090
2. EZW Utility -- Perm.	6.5	2	0.072	-0.054	0.126	0.079	-0.056	0.135
3. EZW Utility -- Temp.	6.5	2	0.047	0.041	0.006	0.047	0.040	0.006
4. Power Utility -- Baseline	4.0	0.25	0.112	0.104	0.008	0.097	0.106	-0.009
5. Power Utility -- Perm./One Period	4.0	0.25	0.066	0.030	0.036	0.073	0.028	0.045
6. EZW Utility -- No Disasters	6.5	2	0.046	0.042	0.005	0.046	0.042	0.005
	10.0	2	0.047	0.040	0.007	0.047	0.040	0.007
7. Power Utility -- No Disasters	4.0	0.25	0.119	0.117	0.003	0.120	0.117	0.003
8. EZW Utility -- Gamma Shocks	4.5	2	0.049	0.035	0.014	0.050	0.037	0.013
	6.5	2	0.055	0.020	0.034	0.056	0.022	0.034
	8.5	2	0.061	-0.002	0.063	0.062	-0.003	0.064

In all cases, the discount factor is $\exp(-0.034)$. For case 1, the model of consumption dynamics is parameterized according to the estimates presented in tables 1 through 4. Cases 2-7 are variations on this parameterization. Case 8 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. "No Disasters" refers to a case in which the probability of a disaster is set to zero. "Full Sample" refers to a long sample with a representative set of disasters. "Conditional on No Disaster" refers to a long sample in which agents expect disasters to occur with their normal frequency but non actually occur. "Baseline" refers to cases with multi-period disasters with partial recovery. "Temp " refers to a case in which the long run shock is set to zero. "Perm" refers to a case in which the mean and standard deviation of the long run shock are set equal to the mean and standard deviation of the short run shock. "One Period" refers to cases where the disaster is made to occur in a single period.

TABLE VII
Long Term Bonds and Parial Default

	(1)	(2)	(3)	(4)	(5)
Coefficient of relative risk aversion	6.5	6.5	7.5	6.5	6.5
Intertemporal elasticity of substitution	2	2	2	2	2
Rate of time preference	0.034	0.034	0.034	0.034	0.034
Dividend growth for perpetuity	0.9	0.9	0.9	0.9	0.9
Probability of partial default on perpetuity	0.0	0.4	0.4	0.4	0.735
Probability of partial default on one period bond	0.0	0.4	0.4	0.0	0.0
Log expected return on equity	0.058	0.058	0.061	0.058	0.058
Log expected return on one period bond	0.009	0.023	0.013	0.009	0.009
Log expected return on perpetuity	-0.023	-0.007	-0.025	-0.007	0.021
Equity premium	0.049	0.034	0.048	0.049	0.049
Term premium	-0.032	-0.031	-0.038	-0.016	0.012
Average duration of perpetuity in normal times	11.4	9.8	11.9	9.8	7.7




FIGURE I

A Partially Permanent Disaster

Note: The figure plots the evolution of consumption and potential consumption during and after a disaster lasting five periods with $\phi = -0.15$ and $\theta = -0.05$ 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 7 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 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 of 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.





FIGURE IV
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 5

France



Korea






Figure 5 (cont.)

Chile



United.States






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 6.5 and an intertemporal elasticity of substitution of 2. The typical disaster is a disaster that lasts 7 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 7: Posterior Distribution of the Equity Premium





Figure 8: Distribution of the Equity Premium in Data Without Disasters






FIGURE IX

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 X

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 4. The typical disaster is a disaster that lasts one period and in which the short run and long run disaster shocks are equal to -0.44. All other shocks are set to zero.