

On the Size Distribution of Macroeconomic Disasters*

Robert J. Barro and Tao Jin
Harvard University

February 2011

Abstract

The coefficient of relative risk aversion is a key parameter for analyses of behavior toward risk, but good estimates of this parameter do not exist. A promising place for reliable estimation is rare macroeconomic disasters, which have a major influence on the equity premium. The premium depends on the probability and size distribution of disasters, gauged by proportionate declines in per capita consumption or GDP. Long-term national-accounts data for 36 countries provide a large sample of disasters of magnitude 10% or more. A power-law density provides a good fit to the size distribution, and the upper-tail exponent, α , is estimated to be around 4. A higher α signifies a thinner tail and, therefore, a lower equity premium, whereas a higher coefficient of relative risk aversion, γ , implies a higher premium. The premium is finite if $\alpha > \gamma$. The observed premium of 5% generates an estimated γ close to 3, with a 95% confidence interval of 2 to 4. The results are robust to uncertainty about the values of the disaster probability and the equity premium and can accommodate seemingly paradoxical situations in which the equity premium may appear to be infinite.

* This research has been supported by the National Science Foundation. We appreciate helpful comments from Xavier Gabaix, Rustam Ibragimov, Chris Sims, Jose Ursua, and Marty Weitzman.

The coefficient of relative risk aversion, γ , is a key parameter for analyses of behavior toward risk, but good estimates of this parameter do not exist. A promising area for reliable estimation is rare macroeconomic disasters, which have a major influence on the equity premium—see Rietz (1988), Barro (2006), and Barro and Ursua (2008). For 17 countries with long-term data on returns on stocks and short-term government bills, the average annual (arithmetic) real rates of return were 0.081 on stocks and 0.008 on bills (Barro and Ursua [2008, Table 5]). Thus, if we approximate the risk-free rate by the average real bill return, the average equity premium was 0.073. An adjustment for leverage in corporate financial structure, using a debt-equity ratio of 0.5, implies that the unlevered equity premium averaged around 0.05.

Previous research (Barro and Ursua [2008]) sought to explain an equity premium of 0.05 in a representative-agent model calibrated to fit the long-term history of macroeconomic disasters for up to 36 countries. One element in the calibration was the disaster probability, p , measured by the frequency of macroeconomic contractions of magnitude 10% or more. Another feature was the size distribution of disasters, gauged by the observed histogram in the range of 10% and above. Given p and the size distribution, a coefficient of relative risk aversion, γ , around 3.5 accorded with the target equity premium.

The present paper shows that the size distribution of macroeconomic disasters can be characterized by a power law in which the upper-tail exponent, α , is the key parameter. This parametric approach generates new estimates of the coefficient of relative risk aversion, γ , needed to match the target equity premium. We argue that the parametric procedure can generate more accurate estimates than the sample-average approach because of selection problems related to missing data for the largest disasters. In addition, confidence sets for the power-law parameters translate into confidence intervals for the estimates of γ .

Section I reviews the determination of the equity premium in a representative-agent model with rare disasters. Section II specifies a familiar, single power law to describe the size distribution of disasters and applies the results to estimate the coefficient of relative risk aversion, γ . Section III generalizes to a double power law to get a better fit to the observed size distribution of disasters. Section IV shows that the results are robust to reasonable variations in the estimated disaster probability, the target equity premium, and the threshold for disasters (set initially at 10%). Section V considers possible paradoxes involving an infinite equity premium. Section VI summarizes the principal findings, with emphasis on the estimates of γ .

I. The Equity Premium in a Model with Rare Disasters

Barro (2009) works out the equity premium in a Lucas (1978)-tree model with rare but large macroeconomic disasters. (Results for the equity premium are similar in a model with a linear, AK, technology, in which saving and investment are endogenous.) In the Lucas-tree setting, (per capita) real GDP, Y_t , and consumption, $C_t=Y_t$, evolve as

$$(1) \quad \log(Y_{t+1}) = \log(Y_t) + g + u_{t+1} + v_{t+1}.$$

The parameter $g \geq 0$ is a constant that reflects exogenous productivity growth. The random term u_{t+1} , which is i.i.d. normal with mean 0 and variance σ^2 , reflects “normal” economic fluctuations. The random term v_{t+1} picks up low-probability disasters, as in Rietz (1988) and Barro (2006). In these rare events, output and consumption jump down sharply. The probability of a disaster is the constant $p \geq 0$ per unit of time. In a disaster, output contracts by the fraction b , where $0 < b \leq 1$. The distribution of v_{t+1} is given by

$$\begin{aligned} \text{probability } 1-p: & \quad v_{t+1} = 0, \\ \text{probability } p: & \quad v_{t+1} = \log(1-b). \end{aligned}$$

The disaster size, b , follows some probability density function. In previous research, the density for b was gauged by the observed histogram. The present analysis specifies the form of this distribution—as a power law—and estimates the parameters, including the exponent of the upper tail. Note that the expected growth rate, g^* , of consumption and GDP is

$$g^* = g + (1/2) \cdot \sigma^2 - p \cdot E b,$$

where $E b$ is the mean disaster size.

Barro (2009) shows that, with a representative agent with Epstein-Zin (1989)-Weil (1990) preferences, the formula for the unlevered equity premium is, when the period length approaches zero:

$$(2) \quad r^e - r^f = \gamma \sigma^2 + p \cdot E \{ b \cdot [(1-b)^{-\gamma} - 1] \},$$

where r^e is the expected rate of return on unlevered equity (a claim on aggregate consumption flows), r^f is the risk-free rate, and γ is the coefficient of relative risk aversion.¹ The term in curly brackets has a straightforward interpretation under power utility, where γ equals the reciprocal of the intertemporal elasticity of substitution (IES) for consumption. Then this term is the product of the proportionate decline in equity value during a disaster, b , and the excess of marginal utility of consumption in a disaster state compared to that in a normal state, $(1-b)^{-\gamma} - 1$. Note that, in the present setting, the proportionate fall in equity value during a disaster, b , equals the proportionate fall in consumption and GDP during the disaster.

Equation (2) can be expressed as

$$(3) \quad r^e - r^f = \gamma \sigma^2 + p \cdot [E(1-b)^{-\gamma} - E(1-b)^{1-\gamma} - E b].$$

Equation (3) shows that the key properties of the distribution of b are the expectations of the variable $1/(1-b)$ taken to the powers γ and $\gamma-1$. (The $E b$ term has a minor impact.)

¹The present analysis assumes that the representative agent's relative risk aversion is constant. Empirical support for this familiar specification appears in Brunnermeier and Nagel (2008) and Chiappori and Paiella (2008).

Barro and Ursua (2008) studied macroeconomic disasters by using long-term annual data for real per capita consumer expenditure, C , for 24 countries and real per capita GDP (henceforth, called GDP) for 36 countries.² These data go back at least to 1914 and as far back as 1870, depending on availability, and end in 2006. The annual time series, including sources, are available at www.economics.harvard.edu/faculty/barro/data_sets_barro.

Barro and Ursua (2008) followed Barro (2006) by using an NBER (National Bureau of Economic Research)-style peak-to-trough measurement of the sizes of macroeconomic contractions. Starting from the annual time series, proportionate contractions in C and GDP were computed from peak to trough over periods of one or more years, and declines by 10% or greater were considered. This method yielded 99 disasters for C (for 24 countries) and 157 for GDP (36 countries). The average disaster sizes, subject to the threshold of 10%, were similar for the two measures: 0.215 for C and 0.204 for GDP . The mean durations of the disasters were also similar: 3.6 years for C and 3.5 years for GDP . The listing of the disaster events—by country, timing, and size—are in Barro and Ursua (2008, Tables C1 and C2).³

Equation (1) is best viewed as applying to short periods, approximating continuous time. In this setting, disasters arise as downward jumps at an instant of time, and the disaster size, b , has no time units. In contrast, the underlying data on C and GDP are annual flows. In relating the data to the theory, there is no reason to identify disaster sizes, b , with large contractions in C or GDP observed particularly from one year to the next. In fact, the major disaster events—exemplified by the world wars and the Great Depression—clearly feature cumulative declines

²This approach assumes that the same process for generating macroeconomic disasters and the same model of household risk aversion applies to all countries at all points in time. In general, reliable estimation of parameters for a rare-disasters model requires a lot of data coming from a population that can be viewed as reasonably homogeneous. However, Barro and Ursua (2008) found that results on the determinants of the equity premium were similar if the sample were limited to OECD countries.

³These data on disaster sizes are the ones used in the current study, except for a few minor corrections. The values used are in an online appendix.

over several years, with durations of varying length. In Barro (2006) and Barro and Ursua (2008), the disaster jump sizes, b , in the continuous-time model—corresponding to equation (3) for the equity premium—were approximated empirically by the peak-to-trough measures of cumulative, proportionate decline. Barro (2006, Section V) showed that this procedure would be reasonably accurate if the true model were one with discrete periods with length corresponding to the duration of disasters (all with the same duration, say 3-1/2 years).

The peak-to-trough method for gauging disaster sizes has a number of shortcomings, addressed in ongoing research by Nakamura, Steinsson, Barro, and Ursua (2011). In this work, the underlying model features a probability per year, p , of entering into a disaster state. (Disaster events are allowed to be correlated across countries, as in the world wars and the Great Depression.) Disasters arise in varying sizes (including occasional bonanzas), and the disaster state persists stochastically. This specification generates frequency distributions for the cumulative size and duration of disasters. In addition, as in Gourio (2008), post-disaster periods can feature recoveries in the form of temporarily high growth rates.⁴

The most important implication of Nakamura, et al. (2011) for the equity premium comes from the recoveries. Since, on average, only half the decline in consumption during a disaster turns out to be permanent, the model's predicted equity premium falls short of the value in equation (3). The other extensions have less influence on the equity premium, although the stochastic duration of disasters matters because of effects on the correlation between consumption growth and stock returns during disasters.

⁴One non-issue (raised by Constantinides [2008, pp. 343-344] and Donaldson and Mehra [2008, p. 84]) is the apparent mismatch between the units for rates of return—per year—and the measurement of disaster sizes by cumulative declines over multiple years (with a mean duration around 3-1/2 years). As already noted, the peak-to-trough measures of macroeconomic decline are approximations to the model's jump declines, which have no time units.

The present analysis uses the peak-to-trough measures of declines in C and GDP to generate an empirical distribution of disaster sizes, b . Figures 1 and 2 show the corresponding histograms for transformed disaster sizes, $1/(1-b)$, for C and GDP, respectively. The findings in Nakamura, et al. (2011) suggest that these measures will be satisfactory for characterizing the distribution of disaster sizes but that some downward adjustment to the equity premium in equation (3) would be appropriate to account particularly for the partly temporary nature of the typical disaster.

As in previous research, the estimated disaster probability, p , equals the ratio of the number of disasters to the number of non-disaster years. This calculation yields $p=0.0380$ per year for C and $p=0.0383$ for GDP. Thus, disasters (macroeconomic contractions of 10% or more) typically occur around 3 times per century. The U.S. experience for C is comparatively mild, featuring only 2 contractions of 10% or more over 137 years—with troughs in 1921 and 1933. However, for GDP, the U.S. data show 5 contractions of 10% or more, with troughs in 1908, 1914, 1921, 1933, and 1947.⁵

Barro and Ursua (2008, Tables 10 and 11) used the observed histograms for disaster sizes from the C and GDP data (Figures 1 and 2) to compute the expectation (that is, the sample average) of the expression in brackets on the right side of equation (3) for alternative coefficients of relative risk aversion, γ . The resulting values were multiplied by the estimated p to calculate the disaster term on the right side of the equation. The other term on the right side, $\gamma\sigma^2$, was computed under the assumption $\sigma=0.02$ per year. However, as in Mehra and Prescott (1985), this term was trivial, compared to the equity premium of around 0.05, for plausible values of γ (even with higher, but still reasonable, values of σ). Hence, the disaster term ended up doing

⁵The 1947 GDP contraction was associated with the demobilization after WWII and did not involve a decline in C. The 1908 and 1914 GDP contractions featured declines in C but not up to the threshold of 10%.

almost all the work in explaining the equity premium. A key finding was that a γ around 3.5 got the model's equity premium into the neighborhood of the target value of 0.05.

II. Single-Power-Law Distribution

We work with the transformed disaster size,

$$z \equiv 1/(1-b),$$

which is the ratio of normal to disaster consumption or GDP. This variable enters into the formula for the equity premium in equation (3). The threshold for b , taken to be 0.095, translates into one for z of $z_0=1.105$. As b approaches 1, z approaches ∞ , a limiting property that accords with the usual setting for a power-law distribution.

We start with a familiar, single power law, which specifies the density function as

$$(4) \quad f(z) = Az^{-(\alpha+1)},$$

for $z \geq z_0$, where $A > 0$ and $\alpha > 0$. The condition that the density integrate to 1 from z_0 to ∞ implies

$$(5) \quad A = \alpha z_0^\alpha.$$

The power-law distribution in equation (4) has been applied widely in physics, economics, computer science, ecology, biology, astronomy, and so on. For a review, see Mitzenmacher (2003a). Gabaix (2009) provides examples of power laws in economics and finance and discusses forces that can generate these laws. The examples include sizes of cities (Gabaix and Ioannides [2004]), stock-market activity (Gabaix, et al. [2003, 2006]), CEO compensation (Gabaix and Landier [2008]), and firm size (Luttmer [2007]). The power-law distribution has been given many names, including heavy-tail distribution, Pareto distribution, Zipfian distribution, and fractal distribution.

Pareto (1897) observed that, for large populations, a graph of the logarithm of the number of incomes above a level x against the logarithm of x yielded points close to a straight line with slope $-\alpha$. This property corresponds to a density proportional to $x^{-(\alpha+1)}$; hence, Pareto's α corresponds to ours in equation (4). The straight-line property in a log-log graph can be used to estimate α , as done by Gabaix and Ibragimov (2011) using least squares. A more common method uses maximum-likelihood estimation (MLE), following Hill (1975). We use MLE in our study.

In some applications, such as the distribution of income, the power law gives a poor fit to the observed frequency data over the whole range but provides a good fit to the upper tail.⁶ In many of these cases, a double power law—with two different exponents over two ranges of z —fits the data well. For uses of this method, see Reed (2003) on the distribution of income and Mitzenmacher (2003b) on computer file sizes. The double power law requires estimation of a cutoff value, δ , for z , above which the upper-tail exponent, α , for the usual power law applies. For expository purposes, we begin with the single power law, but problems in fitting aspects of the data eventually motivate a switch to the richer specification.

The single-power-law density in equations (4) and (5) implies that the equity premium in equation (3) is given by

$$(6) \quad r^e - r^f = \gamma\sigma^2 + p \cdot \left\{ \left(\frac{\alpha}{\alpha-\gamma} \right) z_0^\gamma - \left(\frac{\alpha}{\alpha+1-\gamma} \right) z_0^{\gamma-1} + \left(\frac{\alpha}{\alpha+1} \right) \cdot \left(\frac{1}{z_0} \right) - 1 \right\}$$

if $\alpha > \gamma$. (This formula makes no adjustment for the partially temporary nature of disasters, as described earlier.) For given p and z_0 , the disaster term on the right side involves a race between γ , the coefficient of relative risk aversion, and α , the tail exponent. An increase in γ raises the

⁶ There have been many attempts to explain this Paretian tail behavior, including Champernowne (1953), Mandelbrot (1960), and Reed (2003).

disaster term, but a rise in α implies a thinner tail and, therefore, a smaller disaster term. If $\alpha \leq \gamma$, the tail is sufficiently thick that the equity premium is infinite. This result corresponds to a risk-free rate, r^f , of $-\infty$. We discuss these possibilities later. For now, we assume $\alpha > \gamma$.

We turn now to estimation of the tail exponent, α . When equation (4) applies, the log likelihood for N independent observations on z (all at least as large as the threshold, z_0) is

$$(7) \quad \log(L) = N \cdot [\alpha \cdot \log(z_0) + \log(\alpha)] - (\alpha + 1) \cdot [\log(z_1) + \dots + \log(z_N)],$$

where we used the expression for A from equation (5). The MLE condition for α follows readily as

$$(8) \quad N/\alpha = \log(z_1/z_0) + \dots + \log(z_N/z_0).$$

We obtained standard errors and 95% confidence intervals for the estimate of α from bootstrap methods.⁷

Table I shows that the point estimate of α for the 99 C disasters is 6.27, with a standard error of 0.81 and a 95% confidence interval of (4.96, 8.12). Results for the 157 GDP disasters are similar: the point estimate of α is 6.86, with a standard error of 0.76 and a 95% confidence interval of (5.56, 8.48).

Given an estimate for α —and given $\sigma=0.02$, $z_0=1.105$, and a value for p (0.0380 for C and 0.0383 for GDP)—we need only a value for γ in equation (6) to determine the predicted equity premium, $r^e - r^f$. To put it another way, we can find the value of γ needed to generate $r^e - r^f = 0.05$ for each value of α . (The resulting γ has to satisfy $\gamma < \alpha$ for $r^e - r^f$ to be finite.) In Table I, the point estimate for α of 6.27 from the single power law for the C data requires $\gamma=3.97$. The corresponding standard error for the estimated γ is 0.51, with a 95% confidence interval of (3.13,

⁷See Efron and Tibshirani (1993). We get similar results based on $-2 \cdot \log(\text{likelihood ratio})$ being distributed asymptotically as a chi-squared distribution with one degree of freedom (see Greene [2002]).

5.13). For the GDP data, the point estimate of γ is 4.33, with a standard error of 0.48 and a 95% confidence interval of (3.50, 5.33).

To assess these results, we now evaluate the fit of the single power law. Figure 1 compares the histogram for the C disasters with the frequency distribution implied by the single power law in equations (4) and (5), using $z_0=1.105$ and $\alpha=6.27$ from Table I. An important inference is that the single power law substantially underestimates the frequency of large disasters. Similar results apply for GDP in Figure 2.

The failures in the single power law are clearer in diagrams for cumulative densities. The straight lines in Figures 3 and 4 show, for C and GDP, respectively, fitted logs of probabilities that transformed disaster sizes exceed the values shown on the horizontal axes. The lines connecting the points show logs of normalized ranks of disaster sizes (as in Gabaix and Ibragimov [2011]). If the specified single power law were valid, the two graphs in each figure should be close to each other over the full range of z . However, the figures demonstrate that the single power laws underestimate the probabilities of being far out in the upper tails.

One way to improve the fits is to allow for a smaller tail exponent at high disaster sizes by generalizing to a double power law. This form specifies an upper-tail exponent, α , that applies for z at or above a cutoff value, $\delta \geq z_0$, and a lower-tail exponent, β , that applies below, for $z_0 \leq z < \delta$. This generalization requires estimation of three parameters: the exponents, α and β , and the cutoff, δ . We still treat the threshold, z_0 , as known and equal to 1.105. We should note that the critical parameter for the equity premium is the upper-tail exponent, α . The lower-tail exponent, β , is unimportant; in fact, the distribution need not follow a power law in the lower part. However, we have to specify a reasonable form for the lower portion to estimate the cutoff, δ , which influences the estimate of α .

III. Double-Power-Law Distribution

The double-power-law distribution, with exponents β and α , takes the form:

$$(9) \quad f(z) = \begin{cases} 0 & \text{if } z < z_0, \\ Bz^{-(\beta+1)} & \text{if } z_0 \leq z < \delta, \\ Az^{-(\alpha+1)} & \text{if } \delta \leq z, \end{cases}$$

where $\beta, \alpha > 0$; $A, B > 0$; $z_0 > 0$ is the known threshold; and $\delta \geq z_0$ is the cutoff separating the lower and upper parts of the distribution. The conditions that the density integrate to 1 over $[z_0, \infty)$ and that the densities be equal just to the left and right of δ imply

$$(10) \quad B = A\delta^{\beta-\alpha},$$

$$(11) \quad \frac{1}{A} = \frac{\delta^{\beta-\alpha}}{\beta} (z_0^{-\beta} - \delta^{-\beta}) + \frac{\delta^{-\alpha}}{\alpha}.$$

The single power law in equations (4) and (5) is the special case of equations (9)-(11) when $\beta=\alpha$.

The position of the cutoff, δ , determines the number, K , among the total observations, N , that lie below the cutoff. The remaining $N-K$ observations are at or above the cutoff. Therefore, the log likelihood can be expressed as a generalization of equation (7) as:

$$(12) \quad \log(L) = N \cdot \log(A) + K \cdot (\beta - \alpha) \cdot \log(\delta) - (\beta + 1) \cdot [\log(z_1) + \dots + \log(z_K)] \\ - (\alpha + 1) \cdot [\log(z_{K+1}) + \dots + \log(z_N)],$$

where A satisfies equation (11).

We use maximum likelihood to estimate α , β , and δ . One complication is that small changes in δ cause discrete changes in K when one or more observations lie at the cutoff. These jumps do not translate into jumps in $\log(L)$ because the density is equal just to the left and right of the cutoff. However, jumps arise in the derivatives of $\log(L)$ with respect to the parameters. This issue does not cause problems in finding numerically the values of (α, β, δ) that maximize $\log(L)$ in equation (12). Moreover, we get virtually the same answers if we rely on the first-

order conditions for maximizing $\log(L)$ calculated while ignoring the jump problem for the cutoff. These first-order conditions are generalizations of equation (8).⁸

In Table I, the sections labeled double power show the point estimates of (α, β, δ) for the C and GDP data. We again compute standard errors and 95% confidence intervals using bootstrap methods. A key finding is that the upper-tail exponent, α , is estimated to be much smaller than the lower-tail exponent, β . For example, for C, the estimate of α is 4.16, standard error = 0.87, with a confidence interval of (2.66, 6.14), whereas that for β is 10.10, standard error = 2.40, with a confidence interval of (7.37, 15.17). The estimates reject the hypothesis $\alpha=\beta$ in favor of $\alpha<\beta$ at low p-values (for C and GDP).

Table I shows that the estimated cutoff value, δ , for the C disasters is 1.38—recall that this value corresponds to the transformed disaster size, $z \equiv 1/(1-b)$. The corresponding cutoff for b is 0.275. With this cutoff, 77 of the C crises fall below the cutoff, whereas 22 are above. The corresponding cutoff for b with the GDP crises is 0.320, implying that 136 events fall below the cutoff, whereas 21 are above. Despite the comparatively small number of crises above the cutoffs, we know from previous research (Barro and Ursua [2008, Tables 10 and 11]) that the really large crises have the main influence on the equity premium. That assessment still holds for the present analysis.

Figure 5 compares the histogram for the C disasters with the frequency distribution implied by the double power law in equations (9)-(11), using $z_0=1.105$, $\alpha=4.16$, $\beta=10.10$, and

⁸The expressions are: $\frac{1}{\alpha} = \left(\frac{1}{N-K}\right) \cdot [\log\left(\frac{z_{K+1}}{\delta}\right) + \dots + \log\left(\frac{z_N}{\delta}\right)]$,

$$\alpha \cdot \left[\log\left(\frac{z_{K+1}}{z_0}\right) + \dots + \log\left(\frac{z_N}{z_0}\right)\right] + \beta \cdot \left[\log\left(\frac{z_1}{z_0}\right) + \dots + \log\left(\frac{z_K}{z_0}\right)\right] = N,$$

$$\frac{\delta}{z_0} = \left[\frac{N\alpha + K \cdot (\beta - \alpha)}{\alpha \cdot (N - K)}\right]^{1/\beta}.$$

$\delta=1.38$ from Table I. Unlike the single power law in Figure 1, the double power law accords well with the histogram. Results are similar for the GDP data (not shown). Figures 6 and 7 provide corresponding information for cumulative densities. Compared with the single power laws in Figures 3 and 4, the double power laws accord much better with the upper-tail behavior. The improved fits suggest that the double power law would be superior for estimating the coefficient of relative risk aversion, γ .

With respect to the equity premium, the key difference in Table I between the double and single power laws is the substantially smaller upper-tail exponents, α . Since the estimated α is now close to 4, rather than exceeding 6, the upper tails are much fatter when gauged by the double power laws. These fatter tails mean that a substantially lower coefficient of relative risk aversion, γ , accords with the target equity premium of 0.05.

Equation (3) still determines the equity premium, $r^e - r^f$. For given γ , a specification of (α, β, δ) , along with $z_0=1.105$, determines the relevant moments of the disaster-size distribution. That is, we get a more complicated version of equation (6). (As before, this formulation does not adjust for the partially temporary nature of macroeconomic disasters.) Crucially, a finite $r^e - r^f$ still requires $\alpha > \gamma$. The results determine the estimate of γ that corresponds to those for (α, β, δ) in Table I (still assuming $\sigma=0.02$ and $p=0.0380$ for C and 0.0383 for GDP). This procedure yields point estimates for γ of 3.00 from the C disasters and 2.75 from the GDP disasters.

As before, we use bootstrap methods to determine standard errors and 95% confidence intervals for the estimates of γ . Although the main parameter that matters is the upper-tail exponent, α , we allow also for variations in β and δ . For the C disasters, the estimated γ of 3.00 (Table I) has a standard error of 0.52, with a 95% confidence interval of (2.16, 4.15). For GDP,

the estimate of 2.75 has a standard error of 0.56, with a confidence interval of (2.04, 4.21). Thus, γ is estimated to be close to 3, with a 95% confidence band of roughly 2 to 4.⁹

Because of the fatter upper tails, the estimated γ around 3 is well below the values around 4 estimated from single power laws (Table I). Given the much better fit of the double power law, we concentrate on the estimated γ around 3. As a further comparison, results based on the observed histograms for C and GDP disasters (Barro and Ursua [2008, Tables 10 and 11]) indicated that a γ in the vicinity of 3.5 was needed to generate the target equity premium of 0.05.

The last comparison reflects interesting differences in the two methods: the moments of the size distribution that determine the equity premium in equation (3) can be estimated from a parametric form (such as the double power law) that accords with the observed distribution of disasters sizes or from sample averages of the relevant moments (corresponding to histograms). A disadvantage of the parametric approach is that misspecification of the functional form—particularly for the far upper tails that have few or no observations—may give misleading results. In contrast, sample averages seem to provide unbiased estimates for any underlying functional form. However, the sample-average approach is sensitive to a selection problem, whereby data tend to be missing for the largest disasters (sometimes because governments have collapsed or are fighting wars). This situation must apply to an end-of-world (or, at least, end-of-country) scenario, discussed later, where $b=1$. The tendency for the largest disasters to be missing from the sample means that the sample-average approach tends to under-estimate the

⁹For the threshold corresponding to $b=0.095$, there are 99 C crises, with a disaster probability, p , of 0.0380 per year and an average for b of 0.215. Using $\gamma=3.00$, the average of $(1-b)^{-\gamma}$ is 2.90 and that for $(1-b)^{1-\gamma}$ is 1.87. For $b \geq 0.275$, corresponding to the cutoff, there are 22 C crises, with $p=0.0077$, average for b of 0.417, average for $(1-b)^{-\gamma}$ of 7.12, and average for $(1-b)^{1-\gamma}$ of 3.45. For GDP, with the threshold corresponding to $b=0.095$, there are 157 crises, with $p=0.0383$ and an average for b of 0.204. Using $\gamma=2.75$, the average of $(1-b)^{-\gamma}$ is 2.58 and that for $(1-b)^{1-\gamma}$ is 1.68. For $b \geq 0.320$, corresponding to the cutoff, there are 21 GDP crises, with $p=0.0046$, average for b of 0.473, average for $(1-b)^{-\gamma}$ of 8.43, and average for $(1-b)^{1-\gamma}$ of 3.60.

fatness of the tails, thereby leading to an overstatement of γ .¹⁰ In contrast, the parametric approach (with a valid functional form) may be affected little by missing data in the upper tail. That is, the estimate of the upper-tail exponent, α , is likely to have only a small upward bias due to missing extreme observations, which have to be few in number. This contrast explains why our estimated γ around 3 from the double power laws (Table I) is noticeably smaller than the value around 3.5 generated by the observed histograms.

IV. Variations in Disaster Probability, Target Equity Premium, and Threshold

We consider now whether the results on the estimated coefficient of relative risk aversion, γ , are robust to uncertainty about the disaster probability, p , the target equity premium, $r^e - r^f$, and the threshold, z_0 , for disaster sizes. For p , the estimate came from all the sample data, not just the disasters: p equaled the ratio of the number of disasters (for C or GDP) to the number of non-disaster years in the full sample. Thus, a possible approach to assessing uncertainty about the estimate of p would be to use a model that incorporates all the data, along the lines of Nakamura, et al. (2011). We could also consider a richer setting in which p varies over time, as in Gabaix (2010). We carry out here a more limited analysis that assesses how “reasonable” variations in p influence the estimates of γ .¹¹

Figure 8 gives results for C, and analogous results apply for GDP (not shown). Recall that the baseline value for p of 0.038 led to an estimate for γ of 3.00, with a 95% confidence interval of (2.16, 4.15). Figure 8 shows that lowering p by a full percentage point (to 0.028)

¹⁰The magnitude of this selection problem has diminished with Ursua’s (2010) construction of estimates of GDP and consumer spending for periods, such as the world wars, where standard data were missing. Recent additions to his data set—not included in our current analysis—are Russia, Turkey, and China (for GDP). As an example, the new data imply that the cumulative contraction in Russia from 1913 to 1921 was 62% for GDP and 71% for C.

¹¹For a given set of observed disaster sizes (for C or GDP), differences in p do not affect the maximum-likelihood estimates for the parameters of the power-law distributions. We can think of differences in p as arising from changes in the overall sample size while holding fixed the realizations of the number and sizes of disaster events.

increases the point estimate of γ to 3.2, whereas raising p by a full percentage point (to 0.048) decreases the point estimate of γ to 2.8. Thus, substantial variations in p have only a moderate effect on the estimated γ .

We assumed that the target equity premium was 0.05. More realistically, there is uncertainty about this premium, which can also vary over time and space (due, for example, to shifts in the disaster probability, p). As with our analysis of p , we consider how reasonable variations in the target premium influence the estimated γ . An allowance for a higher target equity premium is also a way to adjust the model to account for the partly temporary nature of macroeconomic disasters. That is, since equation (3) overstates the model's equity premium when the typical disaster is partly temporary (as described before), an increase in the target premium is a way to account for this overstatement.

Equation (3) shows that variations in the equity premium, $r^e - r^f$, on the left side are essentially equivalent, but with the opposite sign, to variations in p on the right side. Therefore, diagrams for estimates of γ versus $r^e - r^f$ look similar to Figure 8, except that the slope is positive. Quantitatively, for the C data, if $r^e - r^f$ were 0.03, rather than 0.05, the point estimate of γ would be 2.6, rather than 3.0. On the other side, if $r^e - r^f$ were 0.07, the point estimate of γ would be 3.2. Results with GDP are similar. Thus, substantial variations in the target equity premium have only a moderate influence on the estimated γ .

The results obtained, thus far, apply for a fixed threshold of $z_0=1.105$, corresponding to proportionate contractions, b , of size 0.095 or greater. This choice of threshold is arbitrary. In fact, our estimation of the cutoff value, δ , for the double power laws in Table I amounts to endogenizing the threshold that applies to the upper tail of the distribution. We were able to estimate δ by MLE because we included in the sample a group of observations that potentially

lay below the cutoff. Similarly, to estimate the threshold, z_0 , we would have to include observations that potentially lie below the threshold. As with estimates of p , this extension requires consideration of all (or at least more of) the sample, not just the disasters.

As in the analysis of disaster probability and target equity premium, we assess the impact of variations in the threshold on the estimated coefficient of relative risk aversion, γ . We consider a substantial increase in the threshold, z_0 , to 1.170, corresponding to $b=0.145$, the value used in Barro (2006). This rise in the threshold implies a corresponding fall in the disaster probability, p (gauged by the ratio of the number of disasters to the number of non-disaster years in the full sample). For the C data, the number of disasters declines from 99 to 62, and p decreases from 0.0380 to 0.0225. For the GDP data, the number of disasters falls from 157 to 91, and p declines from 0.0383 to 0.0209. That is, the probability of a disaster of size 0.145 or more is about 2% per year, corresponding to roughly 2 events per century.

The results in Table II, for which the threshold is $z_0=1.170$, can be compared with those in Table I, where $z_0=1.105$. For the single power laws, the rise in the threshold causes the estimated exponent, α , to adjust toward the value estimated before for the upper part of the double power law (Table I). Since the upper-tail exponents (α) were lower than the lower-tail exponents (β), the estimated α for a single power laws falls when the threshold rises. For the C data, the estimated α decreases from 6.3 in Table I to 5.5 in Table II, and the confidence interval shifts downward accordingly. The reduction in α implies that the estimated γ declines from 4.0 in Table I to 3.7 in Table II, and the confidence interval shifts downward correspondingly.

Results for the single power law for GDP are analogous.¹²

¹²These results apply even though the higher threshold reduces the disaster probability, p . That is, disaster sizes in the range between 0.095 and 0.145 no longer count. As in Barro and Ursua (2008, Tables 10 and 11), the elimination of these comparatively small disasters has only a minor impact on the model's equity premium and,

With a double power law, the change in the threshold has much less impact on the estimated upper-tail exponent, α , which is the key parameter for the estimated γ . For the C data, the rise in the threshold moves the estimated α from 4.16 in Table I to 4.05 in Table II, and the confidence interval changes correspondingly little.¹³ These results imply that the results for γ also change little, going from a point estimate of 3.00 with a confidence interval of (2.16, 4.15) in Table I to 3.00 with an interval of (2.21, 4.29) in Table II. Results for GDP are analogous. We conclude that a substantial increase in the threshold has little effect on the estimated γ .

V. Can the Equity Premium Be Infinite?

Weitzman (2007), building on Geweke (2001), argues that the equity premium can be infinite (and the risk-free rate minus infinite) when the underlying shocks are log-normally distributed with unknown variance. In this context, the frequency distribution for asset pricing is the t-distribution, for which the tails can be sufficiently fat to generate an infinite equity premium. The potential for an infinite equity premium arises also—perhaps more transparently—in our setting based on power laws.

For a single power law, the equity premium, $r^e - r^f$, in equation (6) rises with the coefficient of relative risk aversion, γ , and falls with the tail exponent, α , because a higher α implies a thinner tail. A finite equity premium requires $\alpha > \gamma$, and this condition still applies with a double power law, with α representing the upper-tail exponent. Thus, it is easy to generate an infinite equity premium in the power-law setting. For a given γ , the tail has only to be sufficiently fat; that is, α has to satisfy $\alpha \leq \gamma$.

hence, on the value of γ required to generate the target premium of 0.05. The more important force is the thickening of the upper tail implied by the reduction of the tail exponent, α .

¹³The rise in the threshold widens the confidence interval for the estimated lower-tail exponent, β . As the threshold rises toward the previously estimated cutoff, δ , the lower tail of the distribution becomes increasingly less relevant.

However, we assume that the equity premium, $r^e - r^f$, equals a known (finite) value, 0.05. The important assumption here is not that the premium equals a particular number, but rather that it lies in an interval of something like 0.03 to 0.07 and is surely not infinite. Our estimation, therefore, assigns no weight to combinations of parameters, particularly of α and γ , that generate a counter-factual premium, such as ∞ . For given α (and the other parameters), we pick (i.e. estimate) γ to be such that the premium equals the target, 0.05. Estimates constructed this way always satisfy $\alpha > \gamma$ and, therefore, imply a finite equity premium.

The successful implementation of this procedure depends on having sufficient data so that there are enough realizations of disasters to pin down the upper-tail exponent, α , within a reasonably narrow range. Thus, it is important that the underlying data set is very large in a macroeconomic perspective: 2963 annual observations on consumer expenditure, C, and 4653 on GDP. Consequently, the numbers of disaster realizations—99 for C and 157 for GDP—are sufficient to generate reasonably tight confidence intervals for the estimates of α .

Although our underlying data set is much larger than those usually used to study macroeconomic disasters, even our data cannot rule out the existence of extremely low probability events of astronomical size. Our estimated disaster probabilities, p , were 3.8% per year for C and GDP, and the estimated upper-tail exponents, α , were close to 4 (Table I). Suppose that there were a far smaller probability p^* , where $0 < p^* \ll p$, of experiencing a super disaster; that is, one drawn from a size distribution with a much fatter tail, characterized by an exponent α^* , where $0 < \alpha^* \ll \alpha$. If p^* is extremely low, say 0.01% per year, there is a good chance of seeing no realizations of super disasters even with 5000 observations. Thus, our data cannot rule out the potential for these events, and these far-out possibilities may matter. In particular, regardless of how low p^* is, to fit the target equity premium of 0.05, the coefficient of

relative risk aversion, γ , has to satisfy $\gamma < \alpha^*$ to get a finite equity premium.¹⁴ If α^* can be arbitrarily low (a possibility not ruled out by direct observation when p^* is extremely low), the estimated γ can be arbitrarily close to zero. We, thus, get a reversal of the Mehra-Prescott (1985) puzzle, where the coefficient of relative risk aversion required to match the observed equity premium was excessive by a couple orders of magnitude.¹⁵

Any upper bound $B < 1$ on the potential disaster size, b , would eliminate the possibility of an infinite equity premium. In this sense, the extreme results depend on the possibility of an end-of-the-world event, where $b=1$. To consider this outcome, suppose now that the very small probability p^* refers only to $b=1$. In this case, it is immediate from equation (3) that the equity premium, $r^e - r^f$, is infinite if $\gamma > 0$. Thus, with the assumed form of utility function,¹⁶ any positive probability of apocalypse (which cannot be ruled out by “data”), when combined with an equity premium around 0.05, is inconsistent with a positive degree of risk aversion.

The reference to an end-of-the-world event suggests a possible resolution of the puzzle. The formula for the equity premium in equation (2) involves a comparison of the return on equity, interpreted as a claim on aggregate consumption, with that on a risk-free asset, interpreted as a short-term government bill. However, no claim can deliver risk-free consumption (from whom and to whom?) once the world has ended. Therefore, at least in the limit, we have to allow for risk in the “risk-free” claim.

¹⁴The assumption here, perhaps unreasonable, is that constant relative risk aversion applies arbitrarily far out into the tail of low consumption.

¹⁵This reversal is the counter-part of the one described in Weitzman (2007, p. 1110): “Should we be trying to explain the *puzzle pattern*: why is the actually observed equity premium so embarrassingly *high* while the actually observed riskfree rate is so embarrassingly *low* ... ? Or should we be trying to explain the opposite *antipuzzle pattern*: why is the actually observed equity premium so embarrassingly *low* while the actually observed riskfree rate is so embarrassingly *high* ... ?”

¹⁶The result does not depend on the CRRA form but only on the condition that the marginal utility of consumption approaches infinity as consumption tends toward zero.

Even if we restrict to $b < 1$, a disaster that destroys a large fraction, b , of consumption is likely to generate partial default on normally low-risk assets such as government bills. Empirically, this low return typically does not involve explicit default but rather high inflation and, thereby, low realized real returns on nominally denominated claims during wartime (see Barro [2006, Section I.c]). For the 99 C crises considered in the present analysis, we have data (mainly from Global Financial Data) on real bill returns for 58, of which 33 were during peacetime and 25 involved wars. The median realized real rates of return on bills (arithmetic) were 0.014 in the peacetime crises, similar to that for the full sample, and -0.062 in the wartime crises. Thus, the main evidence for partial default on bills comes from wars that involved macroeconomic depressions.

To generalize the model (without specifically considering war versus peace), suppose that the loss rate on government bills is $\Phi(b)$, where $0 \leq \Phi(b) \leq 1$. We assume $\Phi(0)=0$, so that bills are risk-free in normal times. The formula for the equity premium in equation (2) becomes

$$(13) \quad r^e - r^f = \gamma\sigma^2 + p \cdot E\{[b - \Phi(b)] \cdot [(1 - b)^{-\gamma} - 1]\}.$$

Thus, instead of the loss rate, b , on equity, the formula involves the difference in the loss rates during disasters on equity versus bills, $b - \Phi(b)$. We previously assumed $\Phi(b)=0$, but a more reasonable specification is $\Phi'(b) \geq 0$, with $\Phi(b)$ approaching 1 as b approaches 1. The equity premium in equation (13) will be finite if, as b approaches 1, $b - \Phi(b)$ approaches 0 faster than $(1 - b)^{-\gamma}$ approaches infinity. In particular, the marginal utility of consumption (for a hypothetical survivor) may be infinite if the world ends ($b=1$), but the contribution of this possibility to the equity premium can be nil because no asset can deliver consumption once the world has disappeared.

VI. Summary of Main Findings

The coefficient of relative risk aversion, γ , is a key parameter for analyses of behavior toward risk. We estimated γ by combining information on the probability and sizes of macroeconomic disasters with the observed long-term average equity premium. Specifically, we calculated what γ had to be to accord with a target unlevered equity premium of 5% per year within a representative-agent model that allows for rare disasters.

In our main calibration, based on the long-term global history of macroeconomic disasters, the probability, p , of disaster (defined as a contraction in per capita consumption or GDP by at least 10% over a short period) is 3.8% per year. The size distribution of disasters accords well with a double power law, with an upper-tail exponent, α , of about 4. The resulting estimate of γ is close to 3, with a 95% confidence interval of 2 to 4. This finding is robust to whether we consider consumer expenditure or GDP and to variations in the estimated disaster probability, p , the target equity premium, and the threshold for the size distribution. The results can also accommodate seemingly paradoxical situations in which the equity premium may appear to be infinite.

Table I			
Single and Double Power Laws, Threshold is $z_0=1.105$			
Parameter	Point estimate	Standard error	0.95 Confidence interval
C data (99 disasters)			
Single power law			
α	6.27	0.81	(4.96, 8.12)
γ	3.97	0.51	(3.13, 5.13)
Double power law			
α	4.16	0.87	(2.66, 6.14)
β	10.10	2.40	(7.37, 15.17)
δ	1.38	0.13	(1.24, 1.77)
γ	3.00	0.52	(2.16, 4.15)
GDP data (157 disasters)			
Single power law			
α	6.86	0.76	(5.56, 8.48)
γ	4.33	0.48	(3.50, 5.33)
Double power law			
α	3.53	0.97	(2.39, 6.07)
β	10.51	3.81	(8.67, 20.98)
δ	1.47	0.15	(1.21, 1.69)
γ	2.75	0.56	(2.04, 4.21)

Note: The single power law, given by equations (4) and (5), applies to transformed disaster sizes, $z=1/(1-b)$, where b is the proportionate decline in C (real personal consumer expenditure per capita) or real GDP (per capita). Disasters are at least as large as the threshold, $z_0=1.105$, corresponding to $b \geq 0.095$. The table shows the maximum-likelihood estimate of the tail exponent, α . The standard error and 95% confidence interval come from bootstrap methods. The corresponding estimates of γ , the coefficient of relative risk aversion, come from calculating the values needed to generate an unlevered equity premium of 0.05 in equation (6) (assuming $\sigma=0.02$ and $p=0.0380$ for C and 0.0383 for GDP). For the double power law, given by equations (10)-(12), the table shows the maximum-likelihood estimates of the two exponents, α (above the cutoff) and β (below the cutoff), and the cutoff value, δ . The corresponding estimates of γ come from calculating the values needed to generate an unlevered equity premium of 0.05 in a more complicated version of equation (6).

Table II			
Single and Double Power Laws with Higher Threshold, $z_0=1.170$			
Parameter	Point estimate	Standard error	0.95 Confidence interval
C data (62 disasters)			
Single power law			
α	5.53	0.85	(4.16, 7.61)
γ	3.71	0.53	(2.83, 4.97)
Double power law			
α	4.05	0.87	(2.81, 6.12)
β	11.36	8.27	(6.63, 39.78)
δ	1.37	0.15	(1.21, 1.86)
γ	3.00	0.54	(2.21, 4.29)
GDP data (91 disasters)			
Single power law			
α	5.67	0.81	(4.39, 7.49)
γ	3.86	0.51	(3.03, 4.99)
Double power law			
α	4.77	1.00	(2.42, 6.24)
β	59.22	22.13	(7.90, 76.73)
δ	1.20	0.17	(1.20, 1.75)
γ	3.41	0.60	(2.04, 4.34)

Note: See the notes to Table I. Disasters are now all at least as large as the threshold $z_0=1.170$, corresponding to $b \geq 0.145$. The disaster probability, p , is now 0.0225 for C and 0.0209 for GDP.

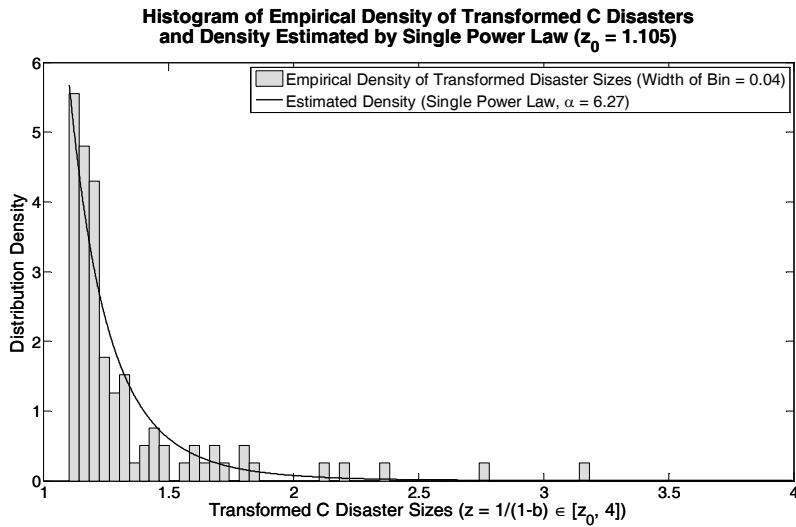


Figure 1

Note: The threshold is $z_0=1.105$, corresponding to $b=0.095$. For the histogram, multiplication of the height on the vertical axis by the bin width of 0.04 gives the fraction of the total observations (99) that fall into the indicated bin. The results for the single power law for C, shown by the curve, correspond to Table I.

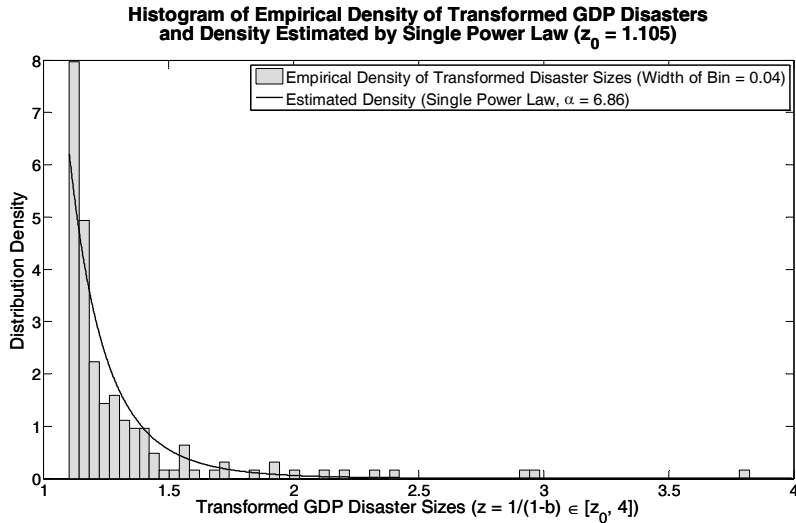
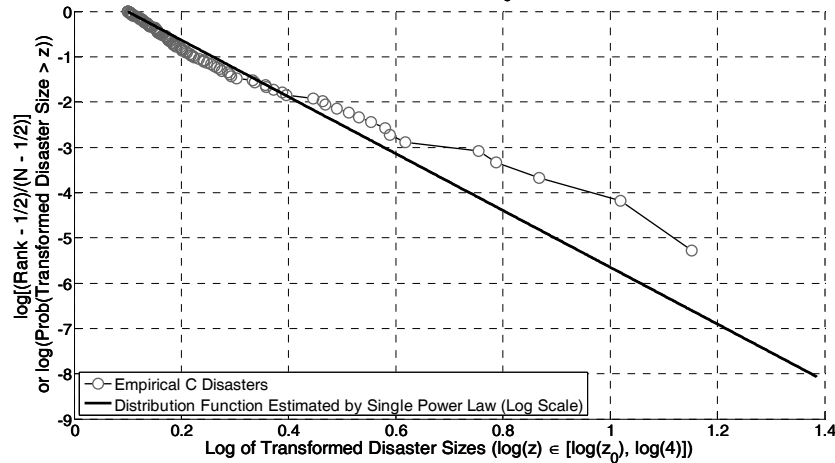


Figure 2

Note: The threshold is $z_0=1.105$, corresponding to $b=0.095$. For the histogram, multiplication of the height on the vertical axis by the bin width of 0.04 gives the fraction of the total observations (157) that fall into the indicated bin. The results for the single power law for GDP, shown by the curve, correspond to Table I.

**Log C Disaster Sizes vs. $\log[(\text{Rank} - 1/2)/(\text{N} - 1/2)]$ and Estimated Distribution Function (Log Scale)
(Single Power Law, $z_0 = 1.105$)**

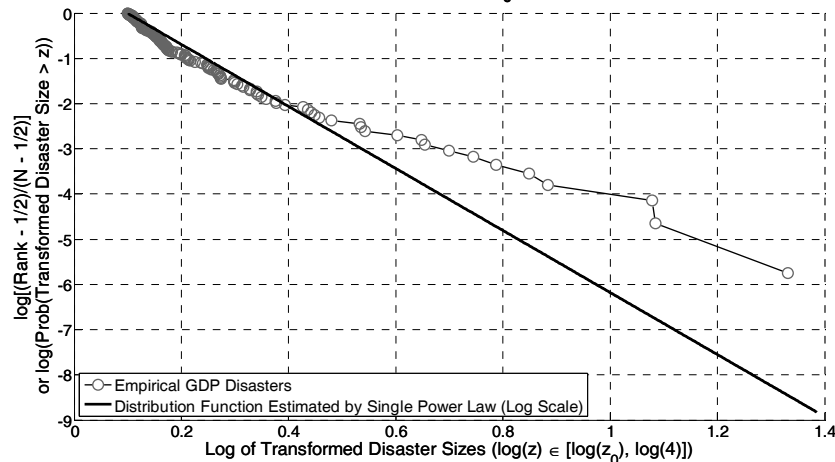


NOTE: $\log[(\text{Rank} - 1/2)/(\text{N} - 1/2)]$ is adopted here instead of $\log(\text{Rank}/\text{N})$. See Gabaix and Ibragimov (2009).

Figure 3

Note: The straight line shows the log of the probability that the transformed disaster size, z , exceeds the quantity shown on the horizontal axis. This line comes from the estimated single power law for C shown in Table I. The lines connecting the points, based on the log of the transformed ranks of the disaster sizes for C, should (if the estimated power law is valid) converge point-wise in probability to the log of the probability given by the straight line (see Gabaix and Ibragimov [2011]).

**Log GDP Disaster Sizes vs. $\log[(\text{Rank} - 1/2)/(\text{N} - 1/2)]$ and Estimated Distribution Function (Log Scale)
(Single Power Law, $z_0 = 1.105$)**



NOTE: $\log[(\text{Rank} - 1/2)/(\text{N} - 1/2)]$ is adopted here instead of $\log(\text{Rank}/\text{N})$. See Gabaix and Ibragimov (2009).

Figure 4

Note: The straight line shows the log of the probability that the transformed disaster size, z , exceeds the quantity shown on the horizontal axis. This line comes from the estimated single power law for GDP shown in Table I. The lines connecting the points, based on the log of the transformed ranks of the disaster sizes for GDP, should (if the estimated power law is valid) converge point-wise in probability to the log of the probability given by the straight line (see Gabaix and Ibragimov [2011]).

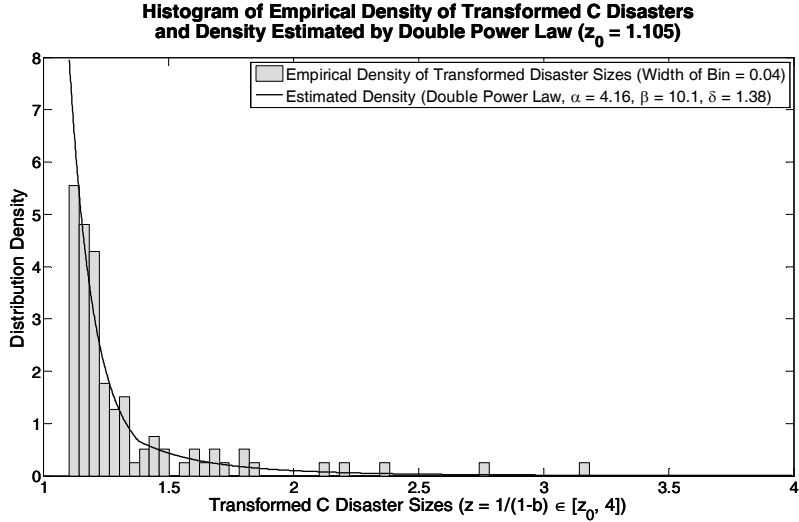
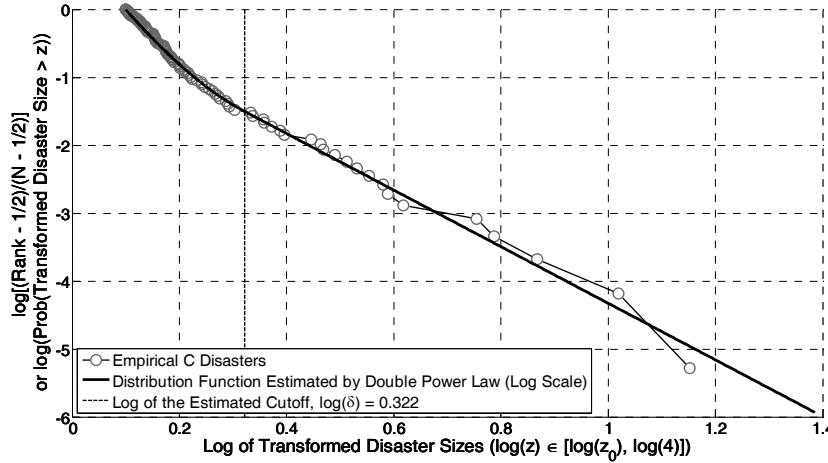


Figure 5

Note: For the histogram, multiplication of the height shown on the vertical axis by the bin width of 0.04 gives the fraction of the total observations (99) that fall into the indicated bin. The results for the double power law for C, shown by the curve, are based on Table I.

Log C Disaster Sizes vs. $\log[(\text{Rank} - 1/2)/(N - 1/2)]$ and Estimated Distribution Function (Log Scale) (Double Power Law, $z_0 = 1.105$)




NOTE: $\log[(\text{Rank} - 1/2)/(N - 1/2)]$ is adopted here instead of $\log(\text{Rank}/N)$. See Gabaix and Ibragimov (2009).

Figure 6

Note: The line with two segments shows the log of the probability that the transformed disaster size, z , exceeds the quantity shown on the horizontal axis. This line comes from the estimated double power law for C in Table I. The lines connecting the points, based on the log of the transformed ranks of the disaster sizes for C, should (if the estimated power law is valid) converge point-wise in probability to the log of the probability given by the line with two segments (see Gabaix and Ibragimov [2011]).

**Log GDP Disaster Sizes vs. $\log[(\text{Rank} - 1/2)/(N - 1/2)]$ and Estimated Distribution Function (Log Scale)
(Double Power Law, $z_0 = 1.105$)**



NOTE: $\log[(\text{Rank} - 1/2)/(N - 1/2)]$ is adopted here instead of $\log(\text{Rank}/N)$. See Gabaix and Ibragimov (2009).

Figure 7

Note: The line with two segments shows the log of the probability that the transformed disaster size, z , exceeds the quantity shown on the horizontal axis. This line comes from the estimated double power law for GDP in Table I. The lines connecting the points, based on the log of the transformed ranks of the disaster sizes for GDP, should (if the estimated power law is valid) converge point-wise in probability to the log of the probability given by the line with two segments (see Gabaix and Ibragimov [2011]).

γ as a Function of p and Lower & Upper Bounds by Percentile Method (γ Corresponds to 5% Equity Premium, Parameters Are Estimated from C Data, $z_0 = 1.105$, $\sigma = 0.020$)

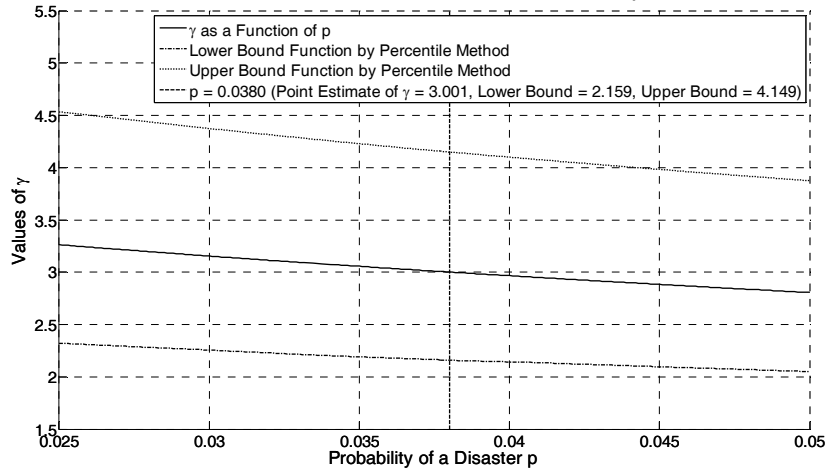


Figure 8

**Estimates of Coefficient of Relative Risk Aversion, γ , for Alternative Disaster Probabilities
Based on C Data with Threshold of $z_0=1.105$**

Note: These results with the higher threshold, $z_0=1.105$, correspond to the estimates for the double power law for C in Table II.

References

- Barro, R.J. (2006): Rare Disasters and Asset Markets in the Twentieth Century, *Quarterly Journal of Economics*, 121, 823–866.
- Barro, R.J. (2009): Rare Disasters, Asset Prices, and Welfare Costs, *American Economic Review*, 99, 243–264.
- Barro, R.J. and J.F. Ursua (2008): Macroeconomic Crises since 1870, *Brookings Papers on Economic Activity*, spring, 255–335.
- Brunnermeier, M.K. and S. Nagel (2008): Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation, *American Economic Review*, 98, 713-736.
- Champernowne, D.G. (1953): A Model of Income Distribution, *Economic Journal*, 63, 318–351.
- Chiappori, P.A. and M. Paiella (2008): Relative Risk Aversion Is Constant: Evidence from Panel Data, Unpublished Manuscript, Columbia University, May.
- Constantinides, G.M. (2008): Comment on Barro and Ursua, *Brookings Papers on Economic Activity*, spring, 341-350.
- Donaldson, J.B. and R. Mehra (2008): Risk-Based Explanations of the Equity Premium, Ch. 2 of J.B. Donaldson and R. Mehra, eds., *Handbook of the Equity Risk Premium*. Amsterdam, the Netherlands: Elsevier.
- Efron, B. and R.J. Tibshirani (1993): *An Introduction to the Bootstrap*. Chapman & Hall/CRC.
- Epstein, L.G. and S.E. Zin. (1989): Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework, *Econometrica*, 57, 937-969.
- Gabaix, X. (2009): Power Laws in Economics and Finance, *Annual Review of Economics*, 1, 255-293.
- Gabaix, X. (2010): Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance, Unpublished Manuscript, New York University, March.

- Gabaix, X., P. Gopikrishnan, V. Plerou, and H.E. Stanley (2003): A Theory of Power Law Distributions in Financial Market Fluctuations, *Nature*, 423, 267–270.
- Gabaix, X., P. Gopikrishnan, V. Plerou, and H.E. Stanley (2006): Institutional Investors and Stock Market Volatility, *Quarterly Journal of Economics*, 121, 461–504.
- Gabaix, X. and R. Ibragimov (2011): Log(RANK-1/2): A Simple Way to Improve the OLS Estimation of Tail Exponents, *Journal of Business Economics and Statistics*, 29, 24-39.
- Gabaix, X. and Y. Ioannides (2004): The Evolution of City Size Distributions, in V. Henderson and J.F. Thisse, eds., *Handbook of Regional and Urban Economics*, v.4. Amsterdam, the Netherlands: North-Holland.
- Gabaix, X. and A. Landier (2008): Why Has CEO Pay Increased so Much? *Quarterly Journal of Economics*, 123, 49-100.
- Geweke, J. (2001): A Note on Some Limitations of CRRA Utility, *Economics Letters*, 71, 341-345.
- Gourio, F. (2008): Disasters and Recoveries, *American Economic Review*, 98, 68-73.
- Greene, W.H. (2002): *Econometric Analysis*, 5th ed. Englewood Cliffs NJ: Prentice Hall.
- Hill, B.M. (1975): A Simple General Approach to Inference about the Tail of a Distribution, *The Annals of Statistics*, 3, 1163-1174.
- Lucas, R.E., Jr. (1978): Asset Prices in an Exchange Economy, *Econometrica*, 46, 1429-1445.
- Luttmer, E.G.J. (2007): Selection, Growth, and the Size Distribution of Firms, *Quarterly Journal of Economics*, 122, 1103-1144.
- Mandelbrot, B. (1960): The Pareto-Levy Law and the Distribution of Income, *International Economic Review*, 1, 79–106.
- Mehra, R. and E.C. Prescott. (1985): The Equity Premium: A Puzzle, *Journal of Monetary Economics*, 15, 145-161.
- Mitzenmacher, M. (2003a): A Brief History of Generative Models for Power Law and Lognormal Distributions, *Internet Mathematics*, 1, 226–251.

- Mitzenmacher, M. (2003b): Dynamic Models for File Sizes and Double Pareto Distributions, *Internet Mathematics*, 1, 305–333.
- Nakamura, E., J. Steinsson, R.J. Barro, and J.F. Ursua (2011): Crises and Recoveries in an Empirical Model of Consumption Disasters, National Bureau of Economic Research, working paper no. 15920.
- Pareto, V. (1897): *Cours d'Economie Politique*, v.2. Paris, France: F. Pichou.
- Reed, W.J. (2003): The Pareto Law of Incomes—an Explanation and an Extension, *Physica A*, 319, 469–486.
- Rietz, T.A. (1988): The Equity Risk Premium: A Solution, *Journal of Monetary Economics*, 22, 117–131.
- Ursua, J.F. (2010): Macroeconomic Archaeology: Unearthing Crises and Long-Term Trends, Unpublished Manuscript, Harvard University, June.
- Weil, P. (1990): Nonexpected Utility in Macroeconomics, *Quarterly Journal of Economics*, 105, 29-42.
- Weitzman, M.L. (2007): Subjective Expectations and Asset-Return Puzzles, *American Economic Review*, 97, 1102-1130.