NBER WORKING PAPER SERIES

ON THE SIZE DISTRIBUTION OF MACROECONOMIC DISASTERS

Robert J. Barro
Tao Jin

Working Paper 15247
http://www.nber.org/papers/w15247

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2009

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 views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

On the Size Distribution of Macroeconomic Disasters
Robert J. Barro and Tao Jin
NBER Working Paper No. 15247
August 2009, Revised September 2009
JEL No. E20,E32,G01,G12

ABSTRACT

In the rare-disasters setting, a key determinant of the equity premium is the size distribution of macroeconomic disasters, gauged by proportionate declines in per capita consumption or GDP. The long-term national-accounts data for up to 36 countries provide a large sample of disaster events of magnitude 10% or more. For this sample, a power-law density provides a good fit to the distribution of the ratio of normal to disaster consumption or GDP. The key parameter of the size distribution is the upper-tail exponent, , estimated to be near 5, with a 95% confidence interval between 3-1/2 and 7. The equity premium involves a race between and the coefficient of relative risk aversion, . A higher signifies a thinner tail and, therefore, a lower equity premium, whereas a higher implies a higher equity premium. The equity premium is finite if 1>. To accord with the observed average unlevered equity premium of around 5%, we get a point estimate for close to 3, with a 95% confidence interval of roughly 2 to 4.

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

Tao Jin
Department of Economics
Harvard University
Cambridge, MA 02138
tjin@fas.harvard.edu

Recent research builds on the rare-disasters idea of Rietz (1988) to explain the long-term average equity premium (Barro [2006], Barro and Ursua [2008]). For a sample of 17 countries with long-term data on returns on stocks and short-term government bills, the average annual real rates of return were 0.081 on stocks and 0.008 on bills (Barro and Ursua [2008, Table 5]). Thus, if the average real bill return is taken as an approximation to the risk-free rate, the average observed equity premium was 0.073. (Since the real bill return is somewhat risky, the risk-free rate likely lies below the average real rate of return on bills, and the equity premium would be somewhat larger than 0.073.) An adjustment for typical leverage in corporate financial structure (using a debt-equity ratio of 0.5) implied that the corresponding unlevered equity premium was around 0.05.

Previous research (Barro and Ursua [2008]) sought to explain an average unlevered equity premium of 5% within a simple representative-agent model calibrated to accord with the long-term history of macroeconomic disasters observed for up to 36 countries. One element in the calibration was the disaster probability, p, measured by the frequency of macroeconomic contractions with magnitude of 10% or more. Another feature of the calibration was the size distribution of disasters, gauged by the observed histogram in the range of 10% and above. Given the value of p and the size distribution, a coefficient of relative risk aversion, γ, of around 3.5 was needed to accord with the target equity premium.

The present study shows that the size distribution of macroeconomic disasters can be characterized empirically by a power law in which the upper-tail exponent is the key parameter. Estimates of the power-law parameters replace the observed histogram in the calibration of the model. This procedure allows for new estimates of γ; that is, for the coefficient of relative risk aversion needed to match the target equity premium. The confidence intervals for the estimates of γ correspond to the confidence sets for the power-law parameters.

Section I reviews the determination of the equity premium in a simple 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—and, thereby, more reliable estimates of γ. Sections IV-VI show that the results are robust to reasonable variations in the disaster probability, p, the target equity premium, and the assumed threshold size for disasters (set initially at 0.10). Section VII 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 that incorporates rare but large macroeconomic disasters. (Results for the equity premium are similar in a model with a linear [AK] technology—this framework allows for endogenous saving and investment.) In the Lucas-tree setting, real GDP, Y_t, and consumption, $C_t=Y_t$, evolve as

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

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 $0<b<1$.[1] The distribution of v_{t+1} is given by

$$\text{probability } 1\text{-}p\text{:} \quad v_{t+1} = 0,$$

$$\text{probability } p\text{:} \quad v_{t+1} = \log(1-b).$$

The disaster size, b, follows some probability density function. In previous research, this density was gauged by the empirical distribution (histogram) of observed disaster sizes. The present analysis specifies the form of this distribution—as a power law—and estimates the key parameters of this distribution, including especially the exponent of the upper tail.

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):

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

where r^e is the expected rate of return on unlevered equity, r^f is the risk-free rate, E is the expectations operator, and γ is the coefficient of relative risk aversion. 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 (with constant values of g, σ, p, and the distribution of b), the proportionate fall in

[1]In this specification, disasters and normal fluctuations have permanent effects on the level of output. Barro, Nakamura, Steinsson, and Ursua (2009) generalize to allow for recoveries from disasters; that is, for periods of abnormally high growth following large contractions.

equity value during a disaster, b, equals the proportionate fall in GDP and consumption during the disaster.

Equation (2) can also be expressed as

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

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 γ-1. (The Eb term has a minor impact.) Note that 1/(1-b) is the ratio of consumption in a normal state to that in a disaster state.

Barro and Ursua (2008) computed values of b using long-term annual data for real per capita consumer expenditure, C, for 24 countries and real per capita GDP, Y, for 36 countries. Time periods went back at least to 1914 and as far back as 1870, depending on the availability of data. The end date was 2006. These data, including sources, are available at http://www.economics.harvard.edu/faculty/barro/data_sets_barro.

Proportionate contractions were computed from peak to trough over periods of one or more years, and contractions, b, of size 0.10 or greater were considered.[2] This method yielded 95 disasters for C and 152 for Y. We use here an updated version of this calculation to isolate 99 disasters for C (for 24 countries) and 157 for Y (36 countries). Figure 1a shows the resulting histogram for disaster sizes based on C, while Figure 1b shows the comparable diagram for Y. The average disaster sizes, subject to the threshold of 0.10, were similar for the two measures—0.215 for C and 0.204 for Y.

[2]Operationally, the threshold for b was 0.0950; that is, 0.10 to two decimal places.

The disaster probability, p, was computed as the ratio of the number of disasters to the number of non-disaster years. This calculation yields similar estimates for the two cases—p=0.0380 per year for C and 0.0383 for Y. Thus, disasters (macroeconomic contractions of 10% or more) typically occur around three times per century. The U.S. experience for C is comparatively mild, featuring only two contractions of size 0.10 or more over 137 years—with troughs in 1921 and 1933. However, for Y, the U.S. data show five contractions of size 0.10 or more, with troughs in 1908, 1914, 1921, 1933, and 1947.[3]

The earlier procedure used the observed histograms for b from the C and Y data (Figures 1a and 1b) to compute the expectation (that is, the sample average) of the expression in brackets on the right side of equation (3) for various values of the coefficient of relative risk aversion, γ. The resulting values were multiplied by the estimated values of p to calculate the disaster term on the right-hand side of the equation. The other term on the right-hand side, $\gamma\sigma^2$, was computed under the assumption σ=0.02 per year. However, as in Mehra and Prescott (1985), this term is trivial, compared to the observed equity premium of around 0.05, for plausible values of γ (even if we assume higher, but still reasonable, values of σ). Thus, the disaster term ends up doing almost all the work in explaining the equity premium. A key finding in Barro and Ursua (2008, Tables 10 and 11) is that a value for γ around 3.5 was needed to get the model's equity premium into the neighborhood of the target value of 0.05.

One shortcoming of the previous methodology is that it sets to zero the probability of ever seeing a disaster of size greater than the largest one observed within sample. This restriction matters because the chance—even if small—of experiencing an extremely large

[3]The 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 0.10.

disaster has a significant impact on the model's predicted equity premium. That is, the large values of b in the tail of the histograms in Figures 1a and 1b count disproportionately in equations (2) and (3) when γ is well above one. A second point is that the previous analysis was able to calculate an interval of values of γ that would generate a satisfactory equity premium (around 0.05) but could not compute a confidence interval for γ. The present analysis allows for the construction of these confidence intervals.

II. Single Power-Law Distribution

In what follows, we work with the transformed disaster size

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

which is the ratio of normal to disaster consumption or GDP. Note that this variable appears in the formulas for the equity premium in equation (3). The threshold for b of 0.095 (see n.2) translates into a threshold for z of z_0=1.105. As b approaches 1 (its largest possible value), z approaches infinity. This limiting property for z accords with the usual setting for a power-law distribution. The histograms for the transformed disaster size, z, are shown in Figures 2a and 2b for C and GDP, respectively.

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

(5) $$f(z) = Az^{-\alpha},$$

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

(6) $$A = (\alpha-1)\cdot z_0^{\alpha-1} .$$

The key parameter that describes the power-law density is the exponent $\alpha>1$.

The power-law distribution, given in equation (5), has been applied widely in physics, economics, computer science, ecology, biology, astronomy, and so on. For a review, see Mitzenmacher (2003a). Gabaix (2009) presents numerous 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]). Because of its many applications, the power-law distribution has been given many names, including heavy-tail distribution, Pareto distribution, Zipfian distribution, and fractal distribution.

In the late 19th century, the Italian economist Vilfredo Pareto 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 sitting close to a straight line with a slope that we call $-\alpha'$ (Pareto [1897]). This property corresponds to a density proportional to $x^{-\alpha'-1}$; hence, α in equation (5) corresponds to $\alpha'+1$. The straight-line property exhibited in a log-log graph can be used to estimate the exponent α; see, for example, Gabaix and Ibragimov (2009), who use a least-squares approach. A more common method to estimate α is maximum likelihood or MLE—see, for example, Howell (2004) and Bauke (2007). 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.[4] In

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

many of these cases, a double power law—which allows for two different exponents over two ranges of the variable z—fits the data well. For uses of this method, see, for example, Reed (2003) on the distribution of income, Mitzenmacher (2003b) on computer file sizes, and Yamamoto and Ohtsuki (2008) on "aggregation-chipping processes." The double power-law specification requires the estimation of a cutoff value, δ, for the variable z, above which the upper-tail exponent, α, for the usual power law applies. In our rare-disasters context, we ultimately adopt this double power-law approach. However, for expository purposes, we begin with the standard single power law. Problems in fitting aspects of the data eventually motivate us to shift to the richer specification.

The single power-law density in equations (5) and (6) implies that the equity premium in equation (3) is given by

$$r^e - r^f = \gamma\sigma^2 + p \cdot \left\{ \left(\frac{\alpha-1}{\alpha-1-\gamma}\right) z_0^{\gamma} - \left(\frac{\alpha-1}{\alpha-\gamma}\right) z_0^{\gamma-1} + \left(\frac{\alpha-1}{\alpha}\right) \cdot \left(\frac{1}{z_0}\right) - 1 \right\} \tag{7}$$

if $\alpha-1>\gamma$. For given p and z_0, the disaster term on the right-hand side involves a race between γ, the coefficient of relative risk aversion, and α, the tail exponent for the disaster-size density function. An increase in γ raises the disaster term. A rise in α implies a thinner tail and, therefore, a smaller disaster term. If $\alpha-1\leq\gamma$, the tail is sufficiently thick that the equity premium is infinite. This result corresponds to a risk-free rate, r^f, of $-\infty$. The formula in equation (7) applies when $\alpha-1>\gamma$.

The outcome with an infinite equity premium resembles findings highlighted by Weitzman (2007). In his framework, sufficient uncertainty about the parameters of the shock distribution can lead to an infinite equity premium, corresponding to a risk-free rate, r^f, of $-\infty$. However, our assumption is that the left-hand side of equation (7) takes on a known, finite value, which we take to be 0.05. Therefore, the maximum-likelihood estimates will satisfy the condition $\alpha\text{-}1\leq\gamma$, so that $r^e\text{-}r^f$ is finite in the model (as it is in the data). For our purposes, the critical point is not that $r^e\text{-}r^f$ takes on a known, fixed value, such as 0.05. Rather, the important property is that we know that $r^e\text{-}r^f$ is substantially positive—greater than, say, 0.03 for sure—but not astronomical—less than, say, 0.07 for sure.

We turn now to estimation of the tail exponent, α. When equation (5) applies, the log likelihood for N observations on z (all at least as large as the threshold, z_0) is given by

(8) $$\log(L) = N\cdot[(\alpha\text{-}1)\cdot\log(z_0) + \log(\alpha\text{-}1)] - \alpha\cdot[\log(z_1) + \ldots + \log(z_N)],$$

where we used the condition for the constant A given in equation (6). The MLE for α follows readily as

(9) $$1/(\alpha\text{-}1) = \log(z_1/z_0) + \ldots + \log(z_N/z_0).$$

We obtained standard errors and 95% confidence intervals for the estimate of α from the percentile bootstrap method.[5]

Table 1 shows that the point estimate of α for the 99 C disasters is 7.27, with a standard error of 0.81 and a 95% confidence interval of (5.96, 9.12). Results based on the 157 GDP disasters are similar: the point estimate of α is 7.86, with a standard error of 0.76 and a 95% confidence interval of (6.56, 9.48).

Given an estimate for α—and given the values σ=0.02, z_0=1.105, and a value for p (0.0380 for C and 0.0383 for Y)—we need only a value for γ in equation (7) to determine the

[5]See Efron and Tibshirani (1993). Results are similar using the adjusted (BC_a) method proposed by Efron (1987). We also get similar results based on -2·log(likelihood ratio) being distributed asymptotically as a chi-squared distribution with one degree of freedom. See Greene (2002) for a discussion of this likelihood-ratio approach.

model's 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 specified value of α. (Note that the resulting γ has to satisfy the condition γ<α -1—in order for r^e-r^f to be finite in equation [7].) For example, the point estimate for α of 7.27 from the C data requires the value γ=3.97 to generate r^e-r^f = 0.05. Similarly, the standard errors and confidence intervals for α translate into standard errors and confidence intervals for γ.

Table 1 shows the resulting estimates of the coefficient of relative risk aversion, γ. For the C data, the point estimate is 3.97, with a standard error of 0.51 and a 95% confidence interval of (3.13, 5.13). For the GDP data, the point estimate is 4.33, with a standard error of 0.48 and a 95% confidence interval of (3.50, 5.33).

The results from the single power law suggest that the coefficient of relative risk aversion, γ, likely lies between 3 and 5-1/2 in order to accord with the target equity premium of 0.05. By way of comparison, the previous results that used the histograms for C and GDP disasters (Barro and Ursua [2008, Tables 10 and 11]) indicated that a value of γ in the vicinity of 3.5 was needed to generate an equity premium of 0.05. Thus, the single power-law results suggest somewhat higher estimates of γ. This result is contrary to our expectations, because we anticipated that the shift to a parameterized distribution would allow for the possibility of disasters higher than any realized within sample and, thereby, lower the value of γ required to fit the observed average equity premium. We think that this conflict between results and expectations reflects shortcomings in the single power-law specification.

Figure 3a compares the observed histogram for the C disasters with the frequency distribution implied by the single power law in equations (5) and (6), using the value α=7.27

(from Table 1), along with z_0=1.105. Figure 3b gives the parallel information for the GDP disasters, using the value α=7.86 (from Table 1), along with z_0=1.105. The vertical axes in Figures 3a and 3b show the distribution densities. For the empirical histograms, multiplication of the bin heights shown on the vertical axis by the bin width of 0.04 gives the fraction of the total observations that fall into each bin. An important inference from the figures is that the single power laws particularly underestimate the frequency of large disasters.

The failures in the single power laws can be seen more clearly in diagrams that assess cumulative densities. The red lines in Figures 4a and 4b show, for C and GDP, respectively, fitted logs of probabilities based on the single power laws that transformed disaster sizes exceed the values shown on the horizontal axes. The blue lines connecting the points show logs of normalized ranks of disaster sizes (discussed in Gabaix and Ibragimov [2009]). If the specified form of the power law is valid, then the two graphs shown in each figure should be close to each other over the full range of z. However, the figures demonstrate that the single power laws substantially underestimate the probabilities of being far out in the upper tails.

One way to account for these discrepancies is that the tail exponent should be lower at high disaster sizes. A tractable way to accord with this pattern is to generalize to a double power law. This new form specifies an upper-tail exponent, α, that applies for z at or above some cutoff value, $\delta \geq z_0$, and a lower-tail exponent, β, that applies below the cutoff, that is, for $z_0 \leq z < \delta$. This generalization requires the estimation of three parameters: the two exponents, α and β, and the cutoff, δ. We still treat the threshold, z_0, as known and equal to 1.105. We consider later the implications of a change in the threshold.

III. Double Power-Law Distribution

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

$$(10) \qquad f(z) = \begin{cases} 0 & \text{if } z < z_0, \\ Bz^{-\beta} & \text{if } z_0 \le z < \delta, \\ Az^{-\alpha} & \text{if } \delta \le z, \end{cases}$$

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

$$(11) \qquad B = A\delta^{\beta-\alpha}$$

and

$$(12) \qquad \frac{1}{A} = \frac{\delta^{\beta-\alpha}}{\beta-1}(z_0^{1-\beta} - \delta^{1-\beta}) + \frac{\delta^{1-\alpha}}{\alpha-1}.$$

The single power law in equations (5) and 6), corresponding to β=α, is the special case of the double power law in equations (10)-(12) when the cutoff, δ, equals the threshold, z_0.

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 (8) as:

$$(13) \qquad \log(L) = N\cdot\log(A) + K\cdot(\beta-\alpha)\cdot\log(\delta) - \beta\cdot[\log(z_1) + \ldots + \log(z_K)] - \alpha\cdot[\log(z_{K+1}) + \ldots + \log(z_N)],$$

where A satisfies equation (12).

We use maximum likelihood to determine the point estimates of the three parameters α, β, and δ. One complication is that small changes in the cutoff, δ, cause discrete changes in K whenever one or more observations lie just 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 do 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 (13). Moreover, it turns out that we get virtually the same answers if we rely on the first-order conditions for maximizing log(L) that are calculated while ignoring the jump problem related to the cutoff. (In this approach, the number K is treated as fixed in the computation of the first-order conditions for maximizing the likelihood.) These first-order conditions are generalizations of equation (9).[6]

The upper part of Table 2 shows the point estimates of (α, β, δ) for the C data, and the lower part shows the corresponding point estimates for the GDP data. We again compute standard errors and 95% confidence intervals using bootstrap methods.

A key finding for C and GDP in Table 2 is that the upper-tail exponent, α, is estimated to be much smaller than the lower-tail exponent, β. For example, for C, the point estimate of α is 5.16, standard error = 0.87, with a confidence interval of (3.66, 7.14), whereas that for β is 11.10, standard error = 2.40, with a confidence interval of (8.37, 16.17). The estimates reject the hypothesis α=β in favor of α<β at a very low p-value (for both C and GDP).

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

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

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

The point estimate of the 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 the proportionate macroeconomic contraction, 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 that lie above the cutoffs, we know from previous research (Barro and Ursua [2008, Tables 10 and 11]) that these really large crises matter a lot for the equity premium. That assessment still holds for the present analysis.

Figure 5a compares the observed histogram for the C disasters with the frequency distribution implied by the double power law in equations (10)-(12), using the values α=5.16, β=11.10, and δ=1.38 (from Table 2), along with z_0=1.105. Figure 5b provides the analogous information for the GDP disasters, using the values α=4.53, β=11.51, and δ=1.47 (from Table 2), along with z_0=1.105. Unlike for the single power laws, considered in Figures 3a and 3b, the values implied by the double power laws seem to accord well with the observed histograms. Figures 6a and 6b provide corresponding information for cumulative densities. Note that, compared with the fits for the single power laws in Figures 4a and 4b, the double power laws accord much better with the upper-tail behavior. These improved fits for the double power laws suggest that they would provide a better basis for estimating the coefficient of relative risk aversion, γ.

With respect to the equity premium, the key difference between the double power laws in Table 2 and the single power laws in Table 1 is the substantially smaller values of the upper-tail exponents, α. Since the values of α are now close to 5, rather than exceeding 7, the upper tails are much fatter when gauged by the double power laws. These fatter tails mean, in turn, that a substantially lower coefficient of relative risk aversion, γ, is needed to accord with the target equity premium of 0.05.

The formula for the equity premium, r^e-r^f, is again given by equation (3). For a given γ, a specification of the parameters (α, β, δ), along with z_0=1.105, determines the moments of the disaster-size distribution that appear in equation (3). That is, we get a more complicated version of the expression given in equation (7). (A finite value for r^e-r^f still requires α-1>γ.) This result allows us to determine the point estimates of γ that correspond to the point estimates of (α, β, δ) shown in Table 2 (still assuming σ=0.02 and p=0.0380 for C and 0.0383 for GDP). This procedure yields a point estimate for γ of 3.02 from the C disasters and 2.96 from the GDP disasters.

We can also proceed as before to determine standard errors and 95% confidence intervals for the estimates of γ. The main distribution parameter that affects the estimate of γ is the upper-tail exponent, α. However, we allow also for variations in β and δ in determining standard errors and confidence intervals for γ. The result is that, for the C disasters, the point estimate for γ of 3.02 has a standard error of 0.52, with a 95% confidence interval of (2.16, 4.15). For the GDP disasters, the point estimate for γ of 2.96 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.

Because of the fatter upper tails with the double power laws, the estimated values of γ are well below those estimated from the single power laws (Table 1). Moreover, as we anticipated, the estimated values of γ—centered around 3—are noticeably lower than the value 3.5 implied by the observed histograms of disaster sizes (Barro and Ursua [2008, Tables 10 and 11]). The main reason is that the double power law, with an upper-tail exponent, α, around 5, gives some weight to outcomes worse than any experienced in the history that underlies the histograms for C and GDP (Figures 2a and 2b). Although the probabilities of these outcomes are very small, the large size (magnified in impact by a value for γ around 3) leads to a non-negligible effect on the equity premium.

IV. Variations in the Disaster Probability

We could, in principle, consider how uncertainty in the estimates of the disaster probability, p, affect the estimates of the coefficient of relative risk aversion, γ. The estimates of p that we used relied on all the sample data, not just the disaster observations. That is, p equaled the ratio of the number of disasters (for C or GDP) to the number of non-disaster years in the full sample. A possible approach to assessing uncertainty about the estimate of p would be to use a model that incorporates all of the data, not just the disaster realizations, along the lines of Barro, Nakamura, Steinsson, and Ursua (2009). We could also consider a richer model in which p varies over time, along the lines of Gabaix (2008).

We carry out here a more limited analysis that assesses how "reasonable" variations in p influence the point estimates and confidence intervals for γ.[7] Figure 7a gives the results for C, and Figure 7b has the results for GDP. In Figure 7a, the baseline value for p of 0.038 led to a point estimate for γ of 3.00, with a 95% confidence interval of (2.16, 4.15). We can see from the figure that lowering p by a full percentage point (to 0.028) increases the point estimate of γ to

[7]Note that, for 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.

only 3.2, whereas raising p by a full percentage point (to 0.048) decreases the point estimate of γ to only 2.8. Figure 7b shows that similar conclusions emerge from the GDP data. In this case, the point estimate of γ goes from 2.75 at the baseline p of 0.038 to 2.9 when p=0.028 and 2.6 when p=0.048. Thus, allowing for substantial uncertainty in the estimates of p has only a moderate impact on the point estimates and confidence intervals for γ.

V. Variations in the Target Equity Premium

Our analysis assumed a known and fixed value of the target (unlevered) equity premium, 0.05. More realistically, there is uncertainty about the true value of this equity premium and, in a richer model, the premium can vary over time and space (due, for example, to shifts in the disaster probability, p). As with variations in p in the previous section, we consider here how reasonable variations in the target equity premium influence the point estimates and confidence intervals for the coefficient of relative risk aversion, γ.

Equation (3) shows that variations in the equity premium, $r^e – r^f$, on the left-hand side are essentially equivalent, but with the opposite sign, to variations in p on the right-hand side. Therefore, diagrams for estimates of the coefficient of relative risk aversion, γ, versus $r^e – r^f$ look similar to Figures 7a and 7b, except that the slopes are now positive. That is, a higher target equity premium requires a higher γ. Quantitatively, for the C data, if $r^e – r^f$ were 0.04, rather than 0.05, the point estimate of γ would be 2.8, rather than 3.0. On the other side, if $r^e – r^f$ were 0.06, the point estimate of γ would be 3.2. Results with the GDP data are similar. Thus, the overall conclusions are similar to those concerning variations in p.

VI. Alternative Thresholds

The results obtained, thus far, apply for a fixed threshold of z_0=1.105, corresponding to proportionate contractions, b, of size 0.10 or greater. This specification of the threshold is arbitrary. In fact, our estimation of the cutoff value, δ, for the double power law in Table 2 amounts to endogenizing the threshold that applies to the upper-tail part of the distribution. We were able to estimate δ by MLE because we included in the sample a group of observations that potentially fell below the cutoff. Similarly, to estimate the threshold, z_0, we would have to include data that potentially lie below the threshold. As with estimates of p, this extension requires consideration of all (or at least more of) the sample data, not just the disaster observations (with sizes above the designated threshold).

We carry out here a limited analysis that assesses the impact of variations in the threshold on the estimated parameters of the power-law distributions (α for the single power law and α, β, and δ for the double power law). Most importantly, we assess the effects 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 =0.15 (the value used in Barro [2006]).[8] This rise in the threshold implies a corresponding fall in the disaster probability, p (gauged still 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.[9] That is, the probability of a disaster of size 0.15 or more is about 2% per year, corresponding to roughly 2 events per century.

[8]More precisely, we use a value of 0.145, corresponding to 0.15 to two decimal places.
[9]The variations in p and the number of disasters do not precisely correspond because a shift in the threshold affects the number of non-disaster years in the full sample.

For the single power law, the new results in Table 3 (with a threshold of z_0=1.170) can be compared with the initial findings in Table 1 (where z_0=1.105). 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 2). Since the upper-tail exponents (α) were lower than the lower-tail exponents (β), the estimated α for a single power law tends to fall when the threshold rises. For the C data, the point estimate of α decreases from 7.3 in Table 1 to 6.5 in Table 3. For GDP, the estimated α declines from 7.9 in Table 1 to 6.7 in Table 3. In each case, the confidence intervals for α shift downward accordingly.

The reductions in α imply a thickening of the upper tail of the size distribution for disasters and lead, thereby, to lower estimates of the coefficient of relative risk aversion, γ (needed to accord with the target equity premium of 0.05). For the C data, the estimated γ declines from 4.0 in Table 1 to 3.7 in Table 3. Similarly, with the GDP data, the value falls from 4.3 in Table 1 to 3.9 in Table 3. The confidence intervals for γ shift downward correspondingly.[10]

With a double power law, the change in the threshold has less of an impact on the estimates of the upper-tail exponent, α, which is the key parameter that affects the estimates of the coefficient of relative risk aversion, γ. For the C data, the rise in the threshold moves the estimated α from 5.16 in Table 2 to 5.05 in Table 4. The confidence intervals for α change

[10]These results apply even though the higher threshold reduces the disaster probability, p. That is, disaster sizes in the range between 0.10 and 0.15 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 computed equity premium and, 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.

correspondingly little.[11] 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 2 to 3.00 with an interval of (2.21, 4.29) in Table 4. Figure 8a describes how the estimates of γ depend on the disaster probability, p, with the higher threshold for the C data.

For the GDP data, the changes from a rise in the threshold are somewhat harder to interpret. The estimated α goes from 4.53 with a confidence interval of (3.39, 7.07) in Table 2 to 5.77 with an interval of (3.42, 7.24) in Table 4. Thus, the confidence interval hardly changes, but the point estimate shifts substantially.[12] The apparent discrepancies in the results arise because the point estimates are not close to the midpoints of the confidence intervals (falling below the center in the first case and above in the second). The main information is in the confidence intervals, not the point estimates.

Similarly, for γ, the estimate goes from 2.75 with a confidence interval of (2.04, 4.21) in Table 2 to 3.41 with an interval of (2.04, 4.34) in Table 4. Again, the confidence interval barely changes, but the point estimate shifts noticeably. Figure 8b describes how the estimates of γ depend on the disaster probability, p, with the higher threshold for the GDP data.

We think that the important result from the C and GDP data is that the confidence intervals for α and γ change little when the threshold rises substantially. Thus, the main message is that a substantial increase in the threshold has only a minor effect on the principal findings about the degree of risk aversion.

[11]Note, however, that the rise in the threshold dramatically widens the confidence interval for the estimate of the lower-tail exponent, β. As the threshold rises toward the previously estimated cutoff value, δ, the lower tail of the size distribution becomes less relevant.

[12]Again, the confidence interval for the lower-tail exponent, β, widens sharply when the threshold rises.

VII. Summary of Main Findings

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

In our main calibration, based on the long-term global history of macroeconomic depressions, 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 density, with an upper-tail exponent, α, estimated to be around five. The resulting point estimate of the key preference parameter γ is close to three, with a 95% confidence interval ranging from two to four. This basic finding is robust to whether we consider contractions of consumer expenditure or GDP and to substantial variations in the disaster probability, p, the target equity premium (set initially at 0.05), and the threshold value for the size distribution.

Table 1 Single Power Law, Estimates of α and γ Threshold is z_0=1.105			
Parameter	**Point estimate**	**Standard error**	**0.95 Confidence interval**
C data (99 disasters)			
α	7.27	0.81	(5.96, 9.12)
γ	3.97	0.51	(3.13, 5.13)
GDP data (157 disasters)			
α	7.86	0.76	(6.56, 9.48)
γ	4.33	0.48	(3.50, 5.33)

Note: The single power law, given by equations (5) and (6), applies to transformed disaster sizes, $z \equiv 1/(1-b)$, where b is the proportionate decline in C (real personal consumer expenditure per capita) or real GDP (per capita). Disasters are all 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 the percentile bootstrap method. 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 (7) (assuming $\sigma=0.02$ and $p=0.0380$ for C and 0.0383 for GDP).

Table 2 Double Power Law, Estimates of α, β, δ, and γ Threshold is z_0=1.105			
Parameter	**Point estimate**	**Standard error**	**0.95 Confidence interval**
C data (99 disasters)			
α	5.16	0.87	(3.66, 7.14)
β	11.10	2.40	(8.37, 16.17)
δ	1.38	0.13	(1.24, 1.77)
γ	3.00	0.52	(2.16, 4.15)
GDP data (157 disasters)			
α	4.53	0.97	(3.39, 7.07)
β	11.51	3.81	(9.67, 21.98)
δ	1.47	0.15	(1.21, 1.69)
γ	2.75	0.56	(2.04, 4.21)

Note: The double power law, given by equations (10)-(12), applies to transformed disaster sizes, $z \equiv 1/(1-b)$, where b is the proportionate decline in C (real personal consumer expenditure per capita) or real GDP (per capita). Disasters are all at least as large as the threshold, $z_0=1.105$ (corresponding to $b \geq 0.095$). 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 standard errors and 95% confidence intervals come from the percentile bootstrap method. 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 (7) (assuming σ=0.02 and p=0.0380 for C and 0.0383 for GDP).

Table 3 Single Power Law, Estimates of α and γ Threshold is z_0=1.170			
Parameter	**Point estimate**	**Standard error**	**0.95 Confidence interval**
C data (62 disasters)			
α	6.53	0.85	(5.16, 8.61)
γ	3.71	0.53	(2.83, 4.97)
GDP data (91 disasters)			
α	6.67	0.81	(5.39, 8.49)
γ	3.86	0.51	(3.03, 4.99)

Note: See the notes to Table 1. Disasters are now all at least as large as the threshold z_0=1.170, corresponding to b>0.145. The disaster probability, p, is now 0.0225 for C and 0.0209 for GDP.

Table 4 Double Power Law, Estimates of α, β, δ, and γ Threshold is z_0=1.170			
Parameter	**Point estimate**	**Standard error**	**0.95 Confidence interval**
C data (62 disasters)			
α	5.05	0.87	(3.81, 7.12)
β	12.36	8.27	(7.63, 40.78)
δ	1.37	0.15	(1.21, 1.86)
γ	3.00	0.54	(2.21, 4.29)
GDP data (91 disasters)			
α	5.77	1.00	(3.42, 7.24)
β	60.22	22.13	(8.90, 77.73)
δ	1.20	0.17	(1.20, 1.75)
γ	3.41	0.60	(2.04, 4.34)

Note: See the notes to Table 2. Disasters are now all at least as large as the threshold z_0=1.170, corresponding to b>0.145. The disaster probability, p, is now 0.0225 for C and 0.0209 for GDP.

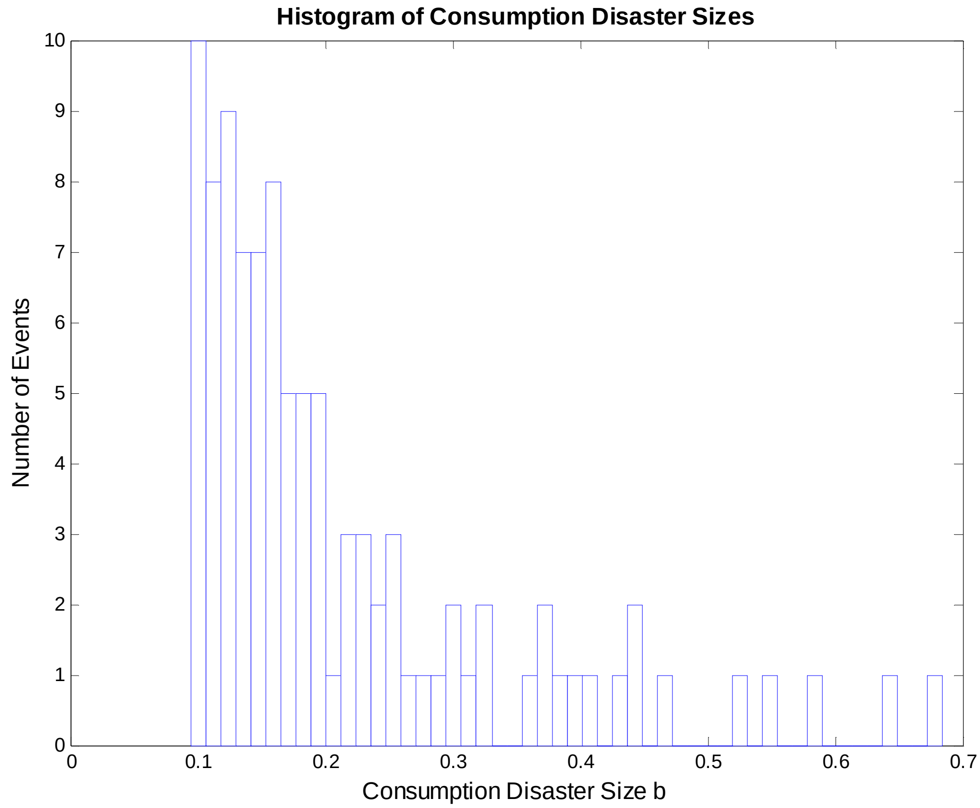

(Total # of Events = 99, Min. = 0.095, Max. = 0.684, Mean = 0.2148, Std. Dev. = 0.1275, # of Bins = 50)

Figure 1a

Note: The threshold is b=0.095.

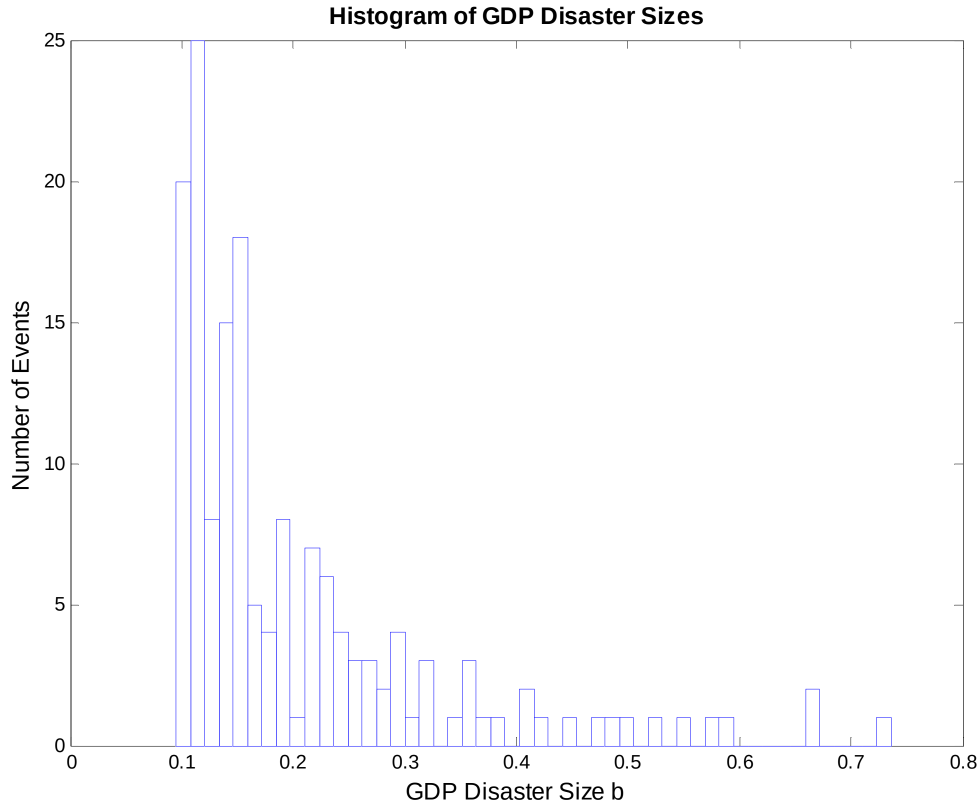

(Total # of Events = 157, Min. = 0.095, Max. = 0.736, Mean = 0.204, Std. Dev. = 0.1264, # of Bins = 50)

Figure 1b

Note: The threshold is b=0.095.

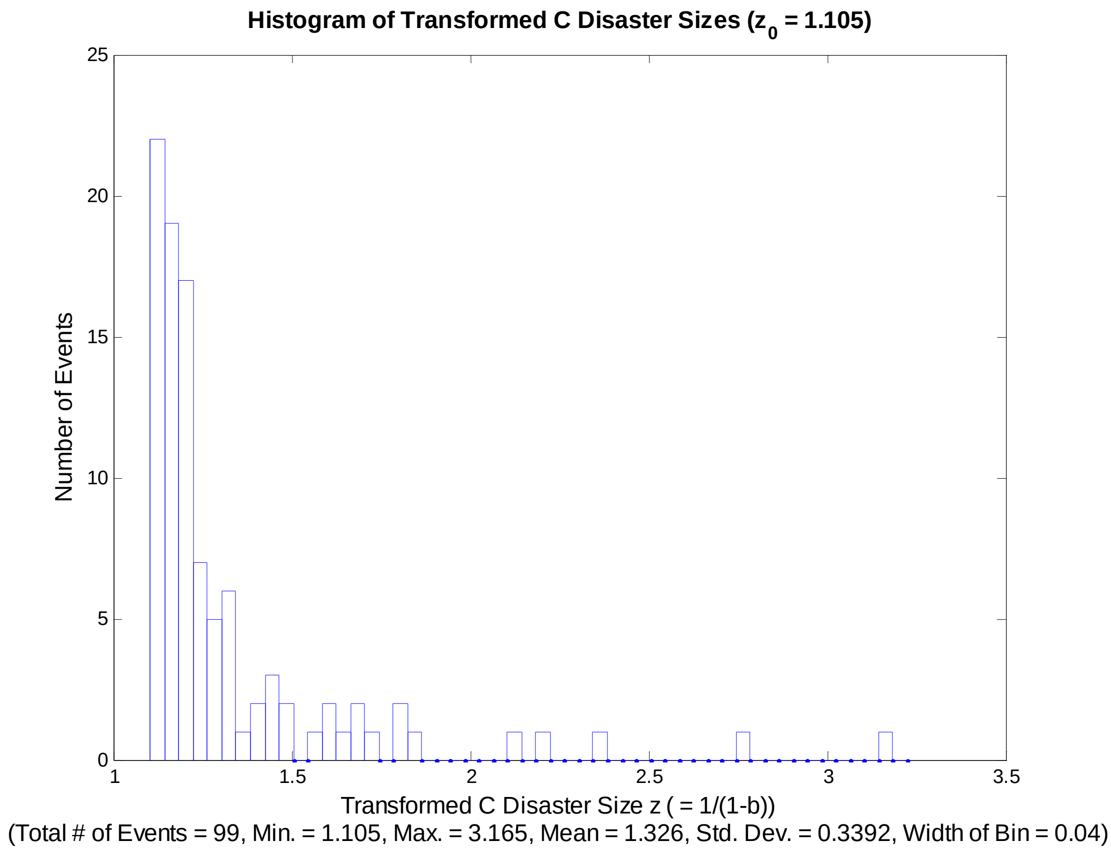

Figure 2a

Note: The threshold is z_0=1.105, corresponding to b=0.095.

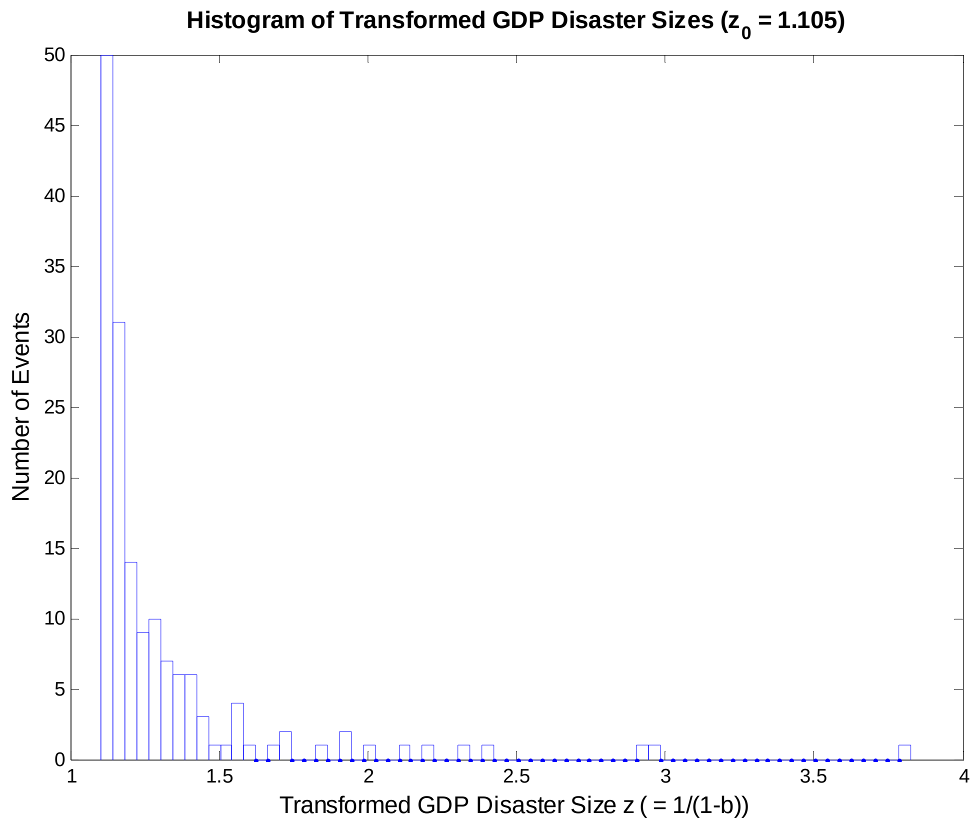

(Total # of Events = 157, Min. = 1.105, Max. = 3.788, Mean = 1.31, Std. Dev. = 0.3613, Width of Bin = 0.04)

Figure 2b

Note: The threshold is z_0=1.105, corresponding to b=0.095.

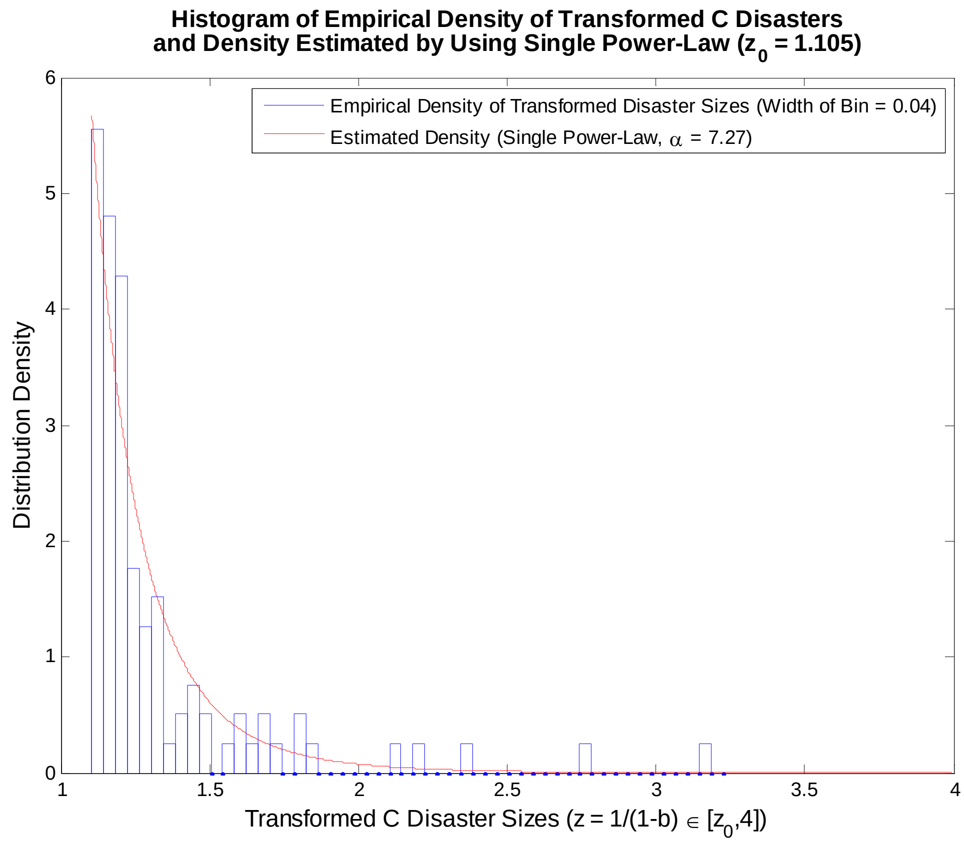

Figure 3a

Note: The results for the single power law for C correspond to Table 1. For the histogram, multiplication of the bin 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.

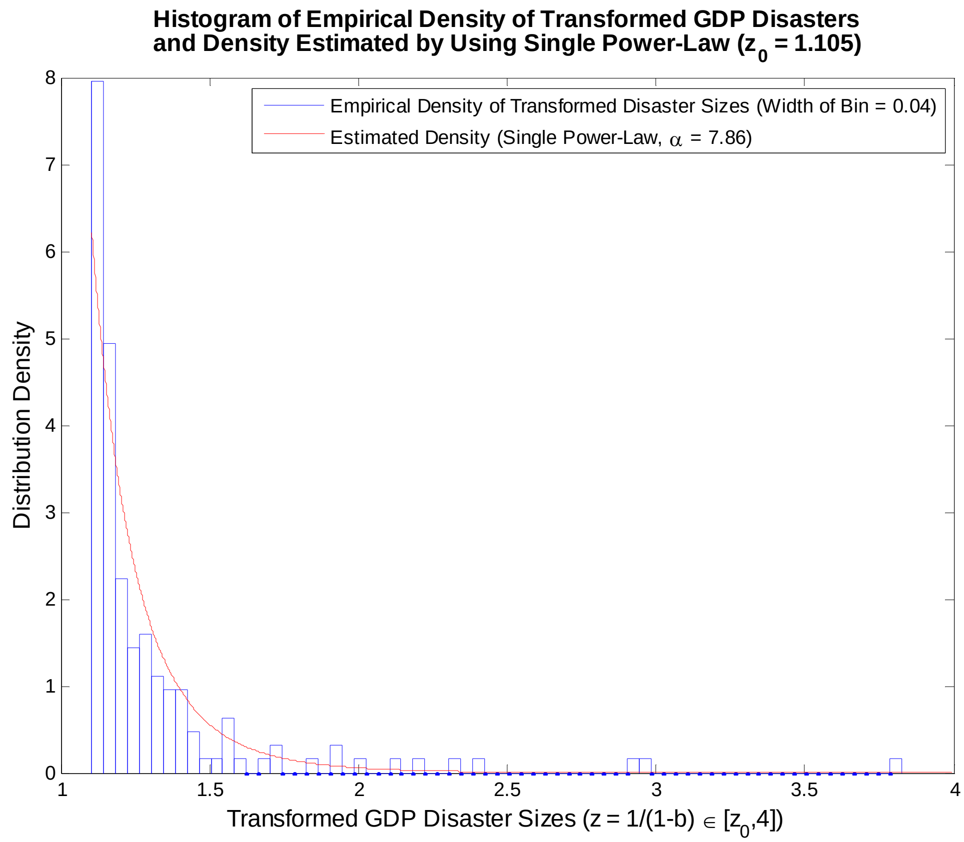

Figure 3b

Note: The results for the single power law for GDP correspond to Table 1. For the histogram, multiplication of the bin height shown 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.

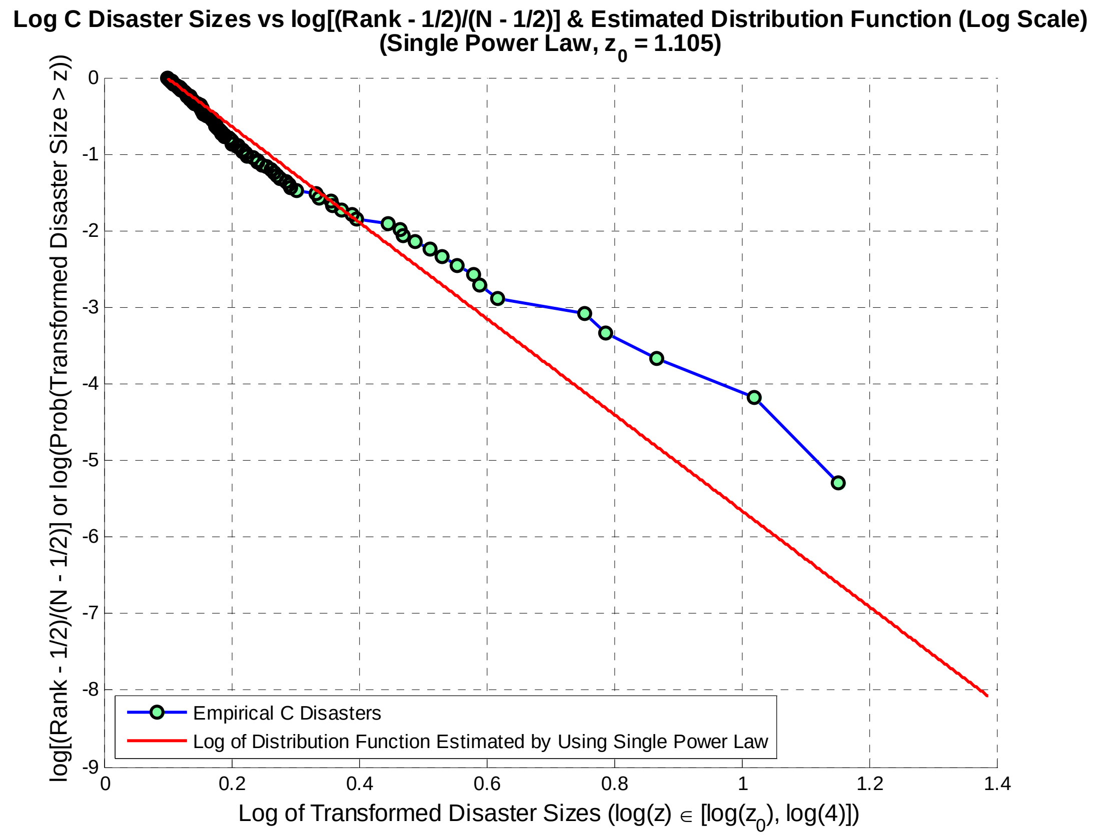

NOTE: log[(Rank - 1/2)/(N - 1/2)] is adopted here instead of log(Rank/N). See Gabaix and Ibragimov (2009).

Figure 4a

Note: The red 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 1. The blue line connecting the points, based on the log of the transformed ranks of the disaster sizes, should (if the estimated power law is valid) converge point-wise in probability to the log of the probability given by the red line (see Gabaix and Ibragimov [2009]).

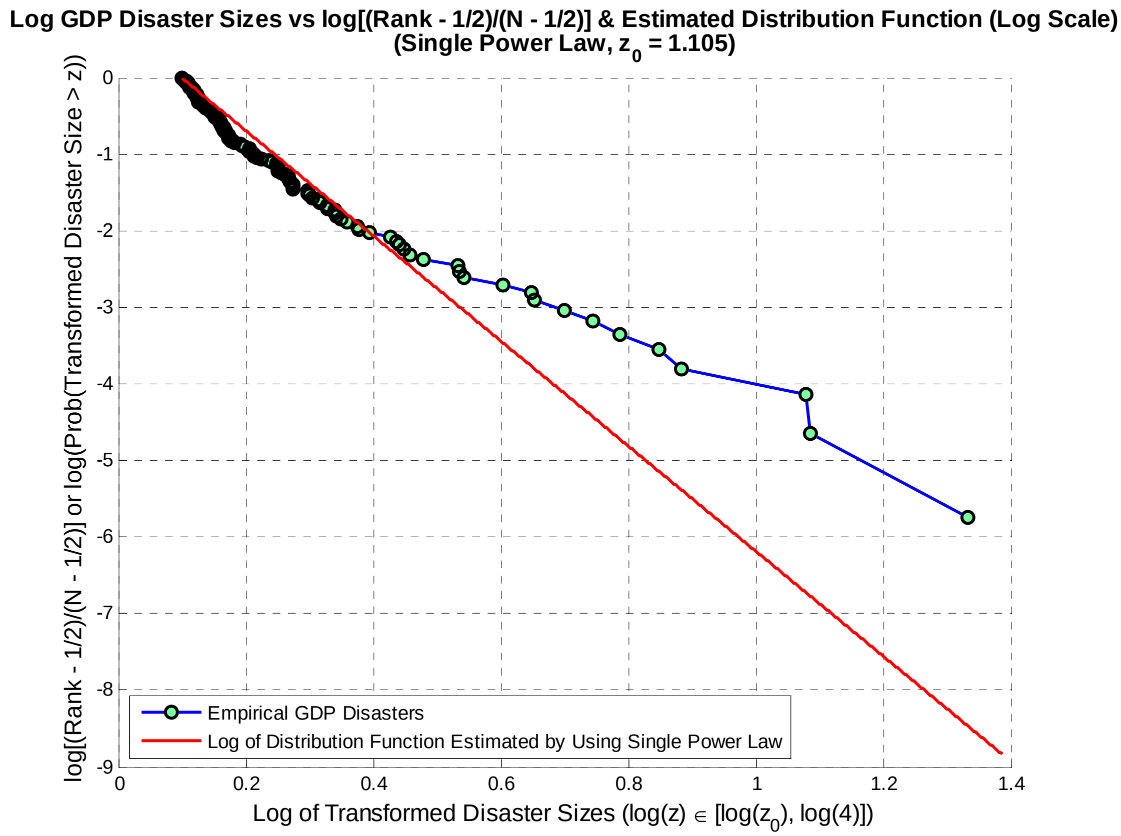

NOTE: log[(Rank - 1/2)/(N - 1/2)] is adopted here instead of log(Rank/N). See Gabaix and Ibragimov (2009).

Figure 4b

Note: The red 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 1. The blue line connecting the points, based on the log of the transformed ranks of the disaster sizes, should (if the estimated power law is valid) converge point-wise in probability to the log of the probability given by the red line (see Gabaix and Ibragimov [2009]).

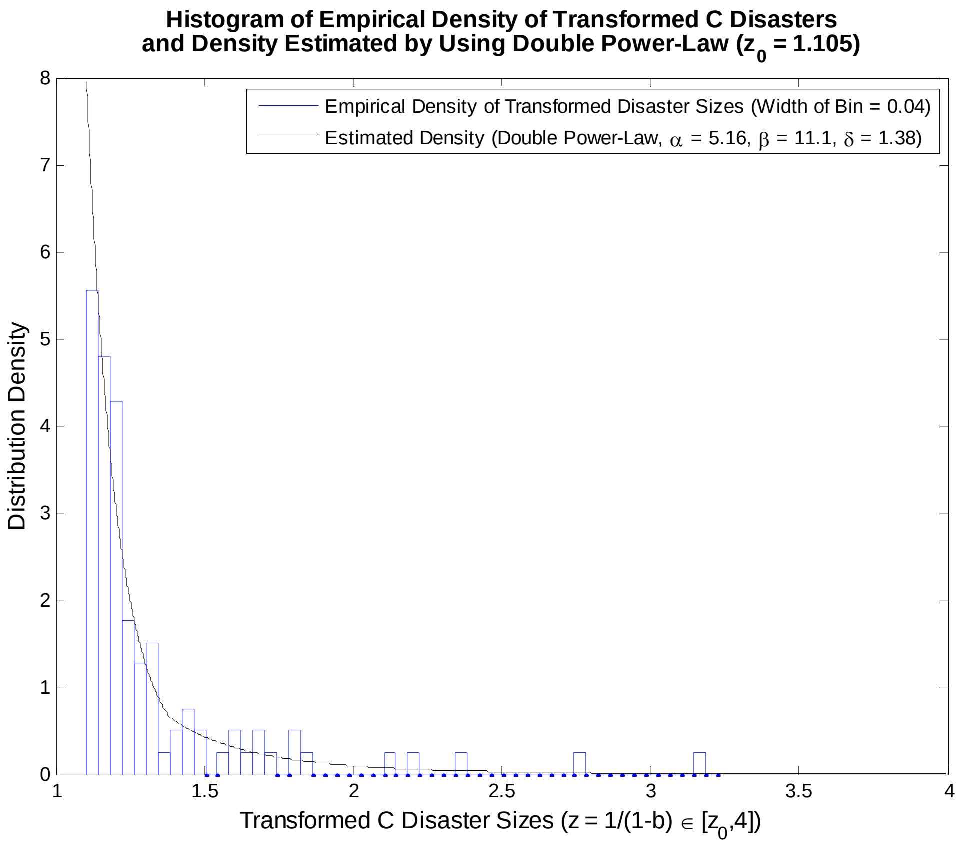

Figure 5a

Note: The results for the double power law for C correspond to Table 2. For the histogram, multiplication of the bin 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.

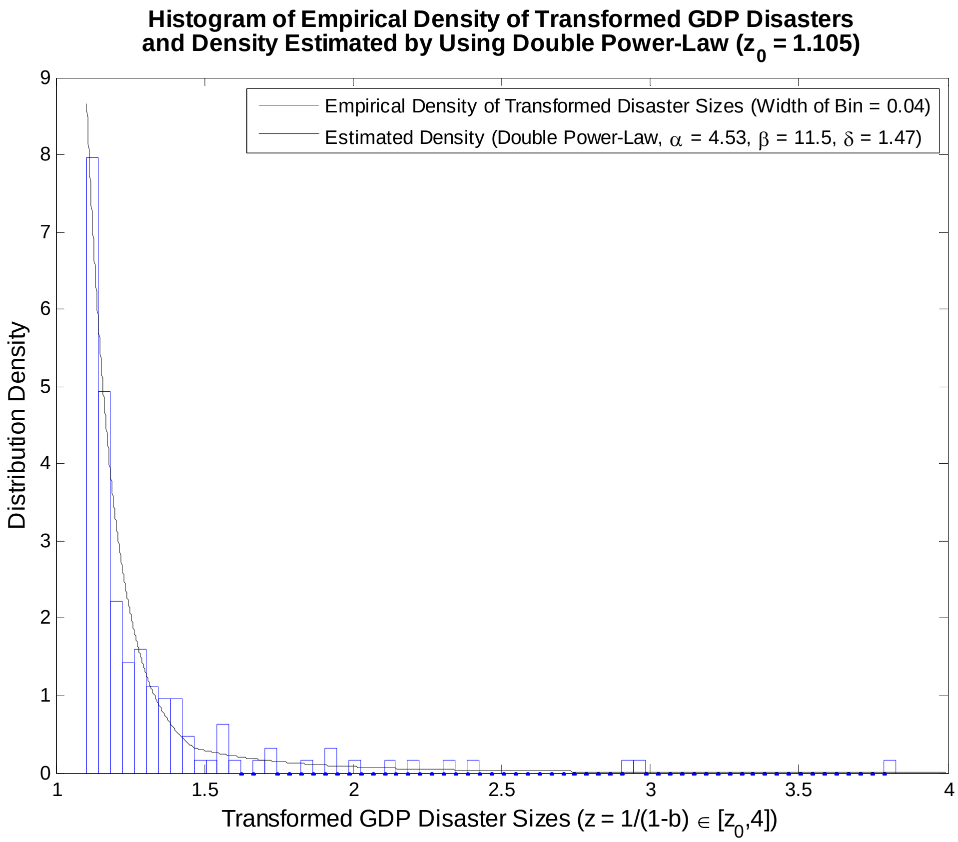

Figure 5b

Note: The results for the double power law for GDP correspond to Table 2. For the histogram, multiplication of the bin height shown 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.

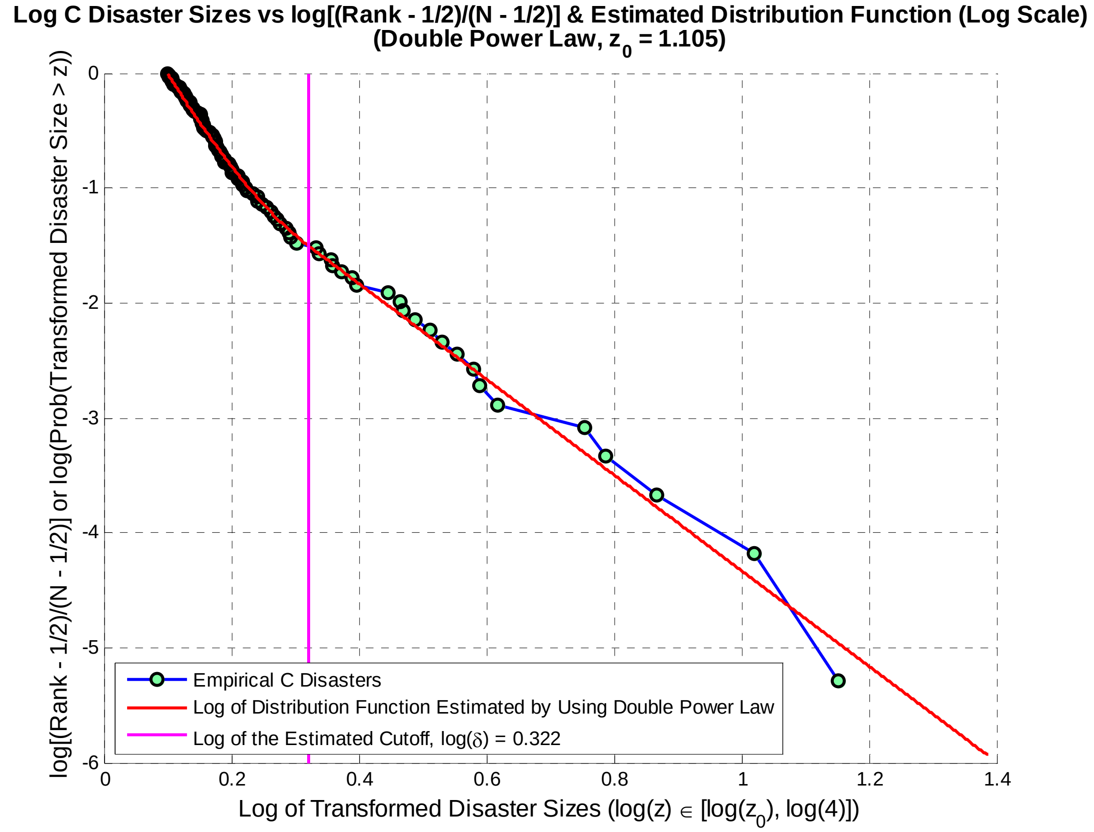

NOTE: log[(Rank - 1/2)/(N - 1/2)] is adopted here instead of log(Rank/N). See Gabaix and Ibragimov (2009).

Figure 6a

Note: The red 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 double power law for C shown in Table 2. The blue line connecting the points, based on the log of the transformed ranks of the disaster sizes, should (if the estimated power law is valid) converge point-wise in probability to the log of the probability given by the red line (see Gabaix and Ibragimov [2009]).

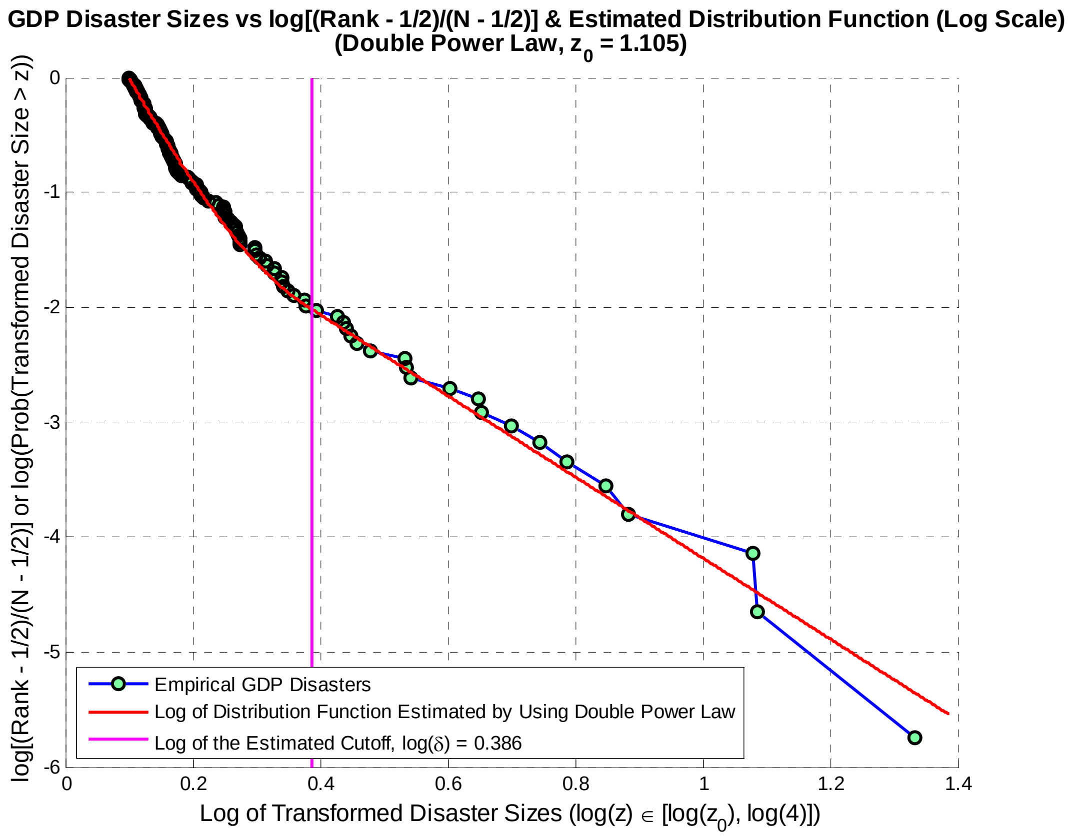

NOTE: log[(Rank - 1/2)/(N - 1/2)] is adopted here instead of log(Rank/N). See Gabaix and Ibragimov (2009).

Figure 6b

Note: The red 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 double power law for GDP shown in Table 1. The blue line connecting the points, based on the log of the transformed ranks of the disaster sizes, should (if the estimated power law is valid) converge point-wise in probability to the log of the probability given by the red line (see Gabaix and Ibragimov [2009]).

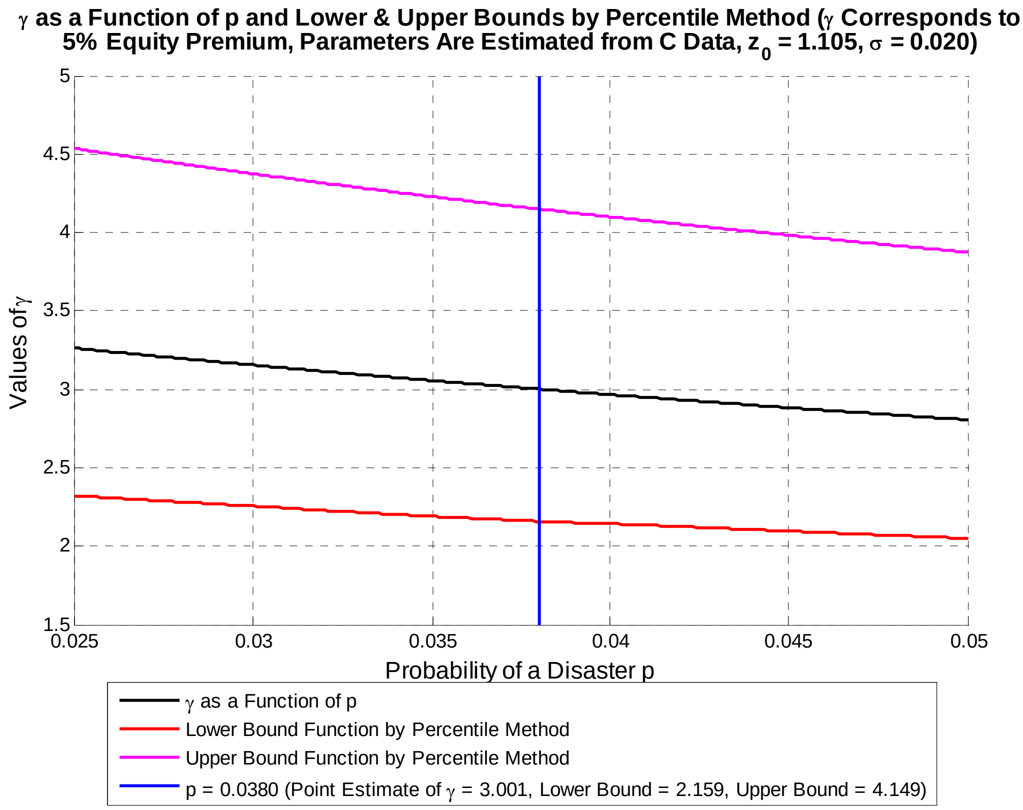

Figure 7a

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 correspond to the estimates for the double power law for C, as described in Table 2.

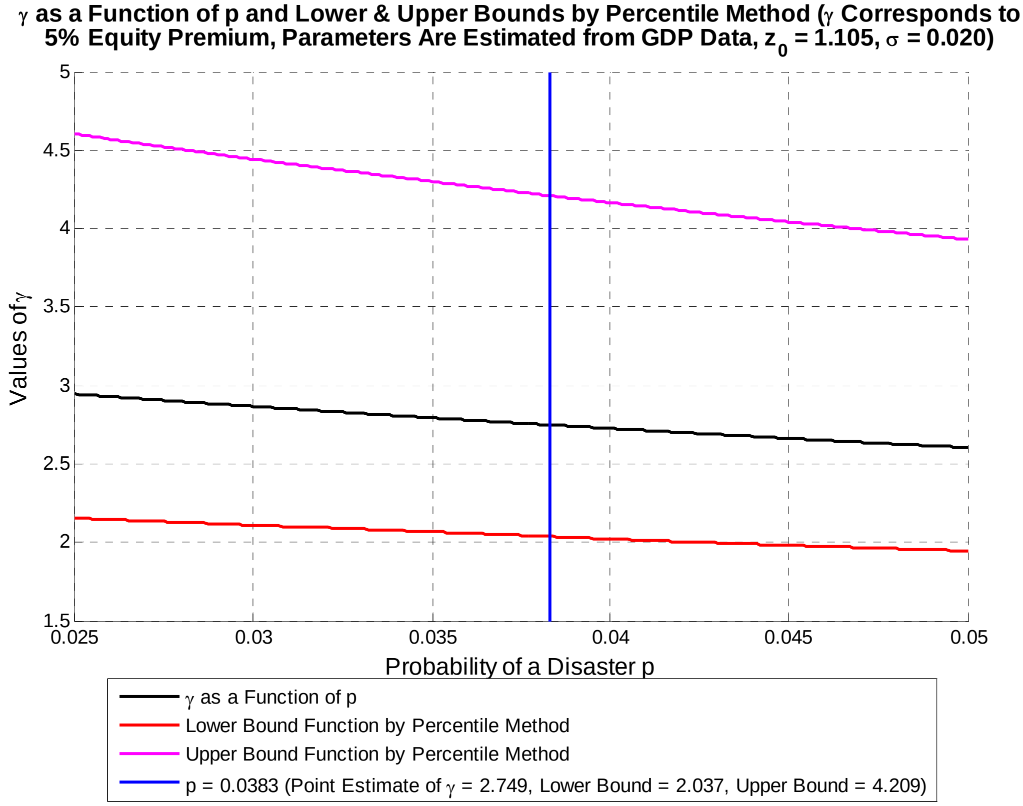

Figure 7b

Estimates of Coefficient of Relative Risk Aversion, γ, for Alternative Disaster Probabilities

Based on GDP Data with Threshold of z_0=1.105

Note: These results correspond to the estimates for the double power law for GDP, as described in Table 2.

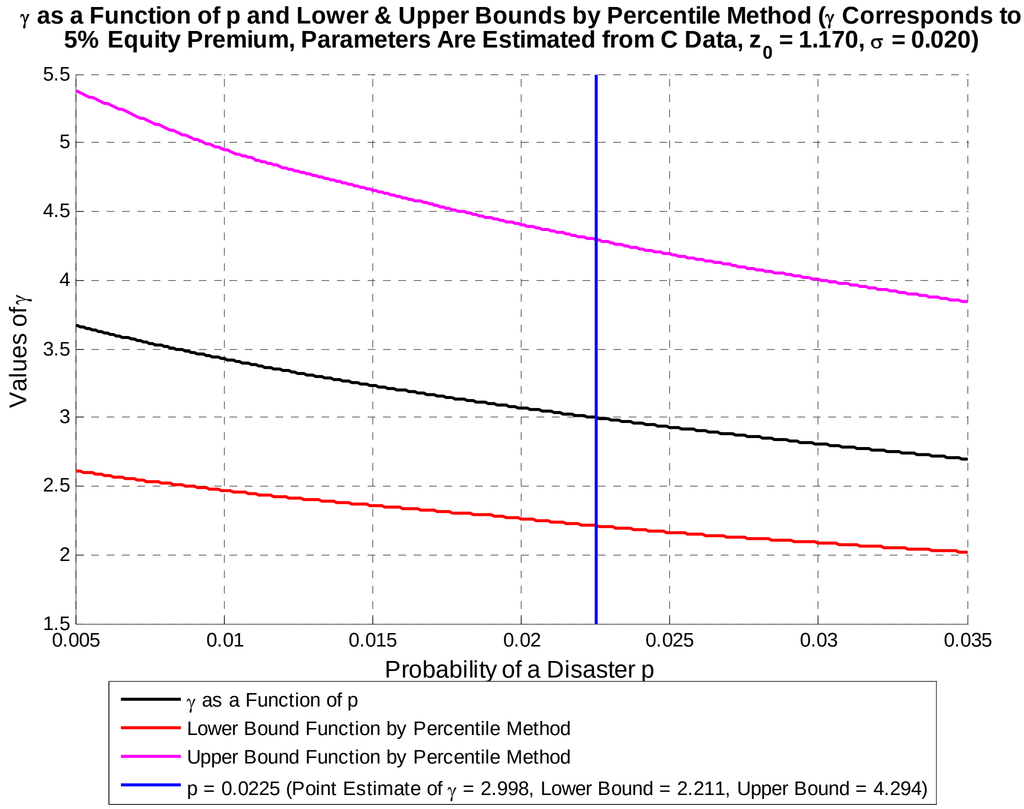

Figure 8a

Estimates of Coefficient of Relative Risk Aversion, γ, for Alternative Disaster Probabilities

Based on C Data with Threshold of z_0=1.170

Note: These results correspond to the estimates for the double power law for C, as described in Table 4.

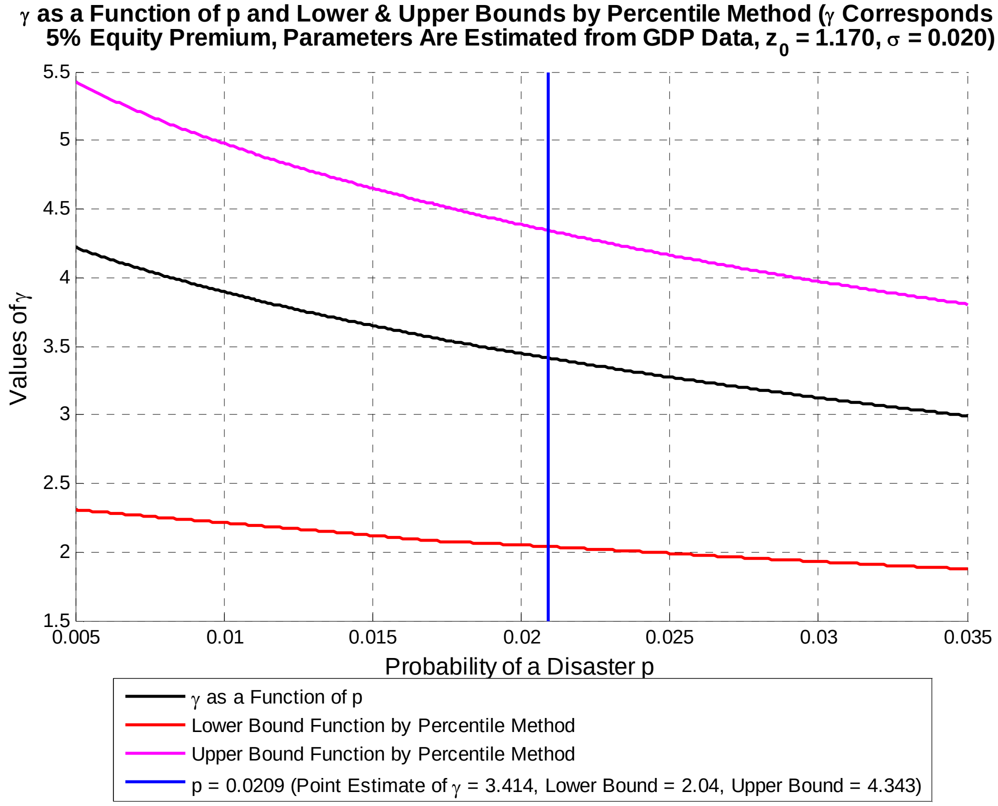

Figure 8b

Estimates of Coefficient of Relative Risk Aversion, γ, for Alternative Disaster Probabilities

Based on GDP Data with Threshold of z_0=1.170

Note: These results correspond to the estimates for the double power law for GDP, as described in Table 4.

References

Barro, R.J. (2006). "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 121, August, 823–866.

Barro, R.J. (2009). "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 99, March, 243–264.

Barro, R.J. and J.F. Ursua (2008). "Consumption Disasters since 1870," *Brookings Papers on Economic Activity*, spring, 255–335.

Barro, R.J., E. Nakamura, J. Steinsson, and J.F. Ursua (2009). "Crises and Recoveries in an Empirical Model of Consumption Disasters," unpublished, Columbia University, July.

Bauke, H. (2007). "Parameter Estimation for Power-Law Distributions by Maximum Likelihood Methods," *European Physical Journal B*, 58, July, 167-173.

Champernowne, D.G. (1953). "A Model of Income Distribution," *Economic Journal,* 63, June, 318–351.

Efron, B. (1987). "Better Bootstrap Confidence Intervals," *Journal of the American Statistical Association*, 82, March, 171–185.

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, July, 937-969.

Gabaix, X. (2008). "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," National Bureau of Economic Research, working paper 13724, January.

Gabaix, X. (2009). "Power Laws in Economics and Finance," *Annual Review of Economics*, 1, January, 255-293.

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, May, 461–504.

Gabaix, X. and R. Ibragimov (2009). "Log(RANK-1/2): A Simple Way to Improve the OLS Estimation of Tail Exponents," forthcoming, *Journal of Business Economics and Statistics*.

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, North-Holland, 2341-2378.

Gabaix, X. and A. Landier (2008). "Why Has CEO Pay Increased so Much?" *Quarterly Journal of Economics*, 123, February, 49-100.

Greene, W.H. (2002). *Econometric Analysis*, 5th ed., Prentice Hall, Englewood Cliffs NJ.

Howell, L.W., Jr. (2004). "Statistical Properties of Maximum Likelihood Estimators of Power Law Spectra Information," *Nuclear Instruments and Methods in Physics Research A,* 521, April, 493-511.

Lucas, R.E., Jr. (1978). "Asset Prices in an Exchange Economy." *Econometrica*, 46, November, 1429-1445.

Luttmer, E.G.J. (2007). "Selection, Growth, and the Size Distribution of Firms," *Quarterly Journal of Economics*, 122, August, 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, March, 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.

Pareto, V. (1897). *Cours d'Economie Politique*, v.2, Paris, F. Pichou.

Reed, W.J. (2003). "The Pareto Law of Incomes—an Explanation and an Extension," *Physica A*, 319, March, 469–486.

Rietz, T.A. (1988). "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 22, July, 117–131.

Weil, P. (1990). "Nonexpected Utility in Macroeconomics," *Quarterly Journal of Economics*, 105, February, 29-42.

Weitzman, M.L. (2007). "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, 97, September, 1102-1130.

Yamamoto, H. and T. Ohtsuki (2008). "Double Power-Law in Aggregation-Chipping Processes," *Physical Review E*, 77, June.