Disasters and Recoveries

François Gourio*

December 2007

1 Introduction

Twenty years ago, Thomas A. Rietz (1988) showed that infrequent, large drops in consumption make the theoretical equity premium large. Recent research has resurrected this 'disaster' explanation of the equity premium puzzle. Robert J. Barro (2006) measures disasters during the XXth century, and finds that they are frequent and large enough, and stock returns low enough relative to bond returns during disasters, to make this explanation quantitatively plausible. Xavier Gabaix (2007) extends the model to incorporate a time-varying incidence of disasters, and he argues that this simple feature can resolve many asset pricing puzzles.

These papers make the simplifying assumption that disasters are permanent. Mathematically, they model log consumption per capita as following a unit root process plus a Poisson jump. However a casual look at the data suggests that disasters are often followed by recoveries. The first contribution of this paper is to measure recoveries and introduce recoveries in the Barro-Rietz model. I find that the effect of recoveries hinges on the intertemporal elasticity of substitution (IES): when the IES is low, recoveries may increase the equity premium implied by the model; but when it is high, the opposite happens.

A second contribution of the paper is to study additional implications of the disaster model.

*I thank Robert Barro, Xavier Gabaix, Ian Martin, Romain Ranciere, Adrien Verdelhan, and participants in a BU macro lunch for discussions or comments. Contact information: Boston University, Department of Economics, 270 Bay State Road, Boston MA 02215. Email: fgourio@bu.edu. Tel.: (617) 353 4534.

Recent empirical research in finance documents that stock returns are forecastable both in the time series and in the cross-section. First, in the time series, a low price-dividend ratio forecasts high excess stock returns and high stock returns. I derive some analytical results which show that it is not easy to replicate these facts in the disaster model, and requires an IES greater than unity. Second, in the cross-section dimension, there is substantial heterogeneity across stocks in their expected (or average) returns. According to the disaster model, the main determinant of an asset's average return should be its exposure to disasters. Looking at the cross-section of stocks, I find only moderate support for this story.

2 Measuring Disasters and Recoveries

Figure 1 plots log GDP per capita for six countries (Germany, Netherlands, the U.S., Chile, Urugay, and Peru). The vertical full lines indicate the start of disasters, and the vertical dashed lines the end of disasters, as defined by Barro.[1]

In many cases, GDP bounces back just after the end of the disaster. These graphs suggest that something is pulling GDP back towards the pre-disaster level. This is, of course, what the neoclassical growth model would predict: following a capital destruction or a temporary decrease in productivity, labor supply and investment are high, and output converges back quickly to its steady-state level.

To quantify the importance of recoveries, Table 1 present some statistics using the entire sample of disasters[2] identified by Barro. Barro measures disasters as the total decline in GDP from peak to through. Using 35 countries, he finds 60 episodes of GDP declines greater than 15% during the XXth century. Because the end of the disaster is the trough, this computation implies that GDP goes up following the disaster. The key question is, How much?

The first column reports the average across countries of the cumulated growth in each of the

[1] The data is from Maddison (2003).

[2] Except the most recent episodes by Argentina, Indonesia and Urugay, for which the next five years of data is not yet available.

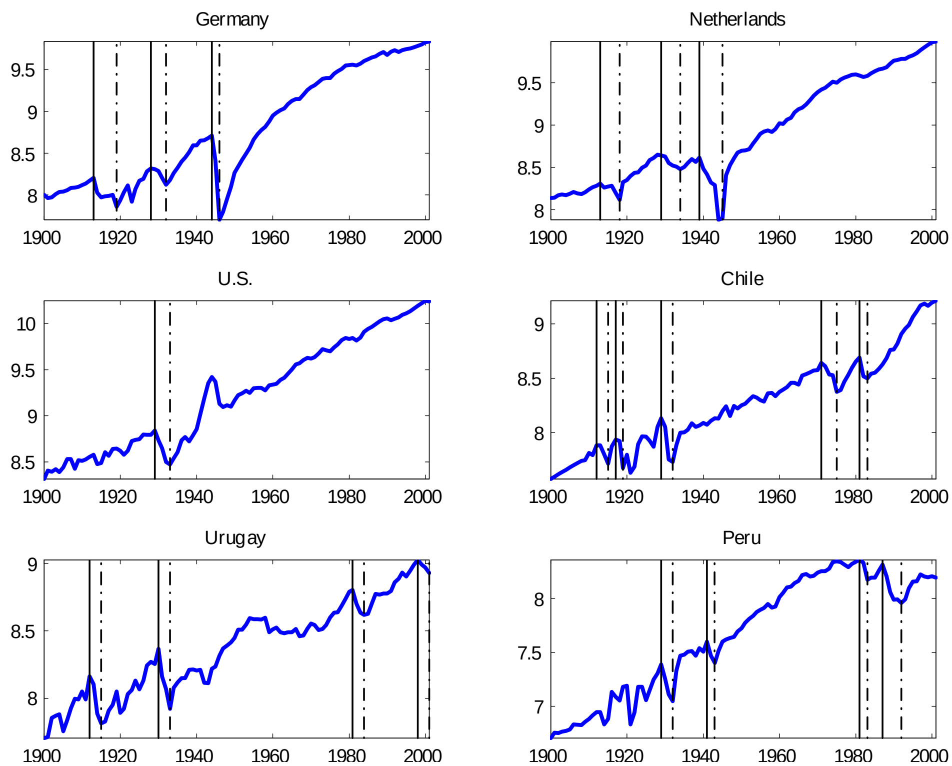

Figure 1: Log GDP per capita (in 1990 dollars) of Germany, Netherlands, the U.S. and Chile, Urugay, and Peru. The disaster start (resp.end) dates are taken from Barro (2006), and are shown with a vertical full (resp. dashed) line.

in %	All disasters (57 events)		Disaster greater than 25% (27 events)	
Years after Trough	Growth from Trough	Loss from previous Peak	Growth from Trough	Loss from previous Peak
0		-29.8		-41.5
1	11.1	-22.8	16.1	-32.7
2	20.9	-16.8	31.3	-24.2
3	26.0	-13.7	38.6	-20.4
4	31.5	-10.2	45.5	-16.9
5	37.7	-6.1	52.2	-13.4

Table 1: Measuring Recoveries. The table reports the average of (a) the growth from the trough to 1,2,3,4,5 years after the trough and (b) the difference from the current level of output to the previous peak level, for 0,1,2,3,4,5 years after the trough.

first five years following a disaster. The average growth rate is 11.1% in the first year after a disaster, and the total growth in the first two years amounts to 20.9%. This is of course much higher than the average growth across these countries over the entire sample, which is just 2.0%. The second column computes how much of the 'gap' from peak to trough is resorbed by this growth, i.e. how much lower is GDP per capita compared to the previous peak. At the trough, on average GDP is 29.8% less than at the previous peak. But on average across countries, this gap is reduced after three years to 13.7%.

Of particular interest are the larger disasters, because diminishing marginal utility implies that agents care enormously about them. Columns 3 and 4 replicate these computations for the subsample of disasters larger than 25%. These disasters are also substantially reversed: a growth of over 30% in the first two years following the disaster nearly erases already half of the decrease in GDP.

There may be better ways to measure recoveries - for instance , I do not take into account trend growth[3] - but I take from these simple computations that at least half of the disaster is, on average, eliminated very quickly.

[3]Many countries have very erratic experiences and it is difficult to define or measure precisely a trend (e.g., because of trend breaks).

3 A Disaster model with Recoveries and Epstein-Zin utility

How do recoveries affect the predictions of the disaster model? This turns out to be more subtle than one might think. To study this question, I extend the Barro-Rietz model and allow for recoveries. For reasons that will become clear, it is useful to introduce Epstein-Zin preferences (1989): utility is defined recursively as

$$V_t = \left(\left(1 - e^{-\rho}\right) C_t^{1-\alpha} + e^{-\rho} E_t \left(V_{t+1}^{1-\theta} \right)^{\frac{1-\alpha}{1-\theta}} \right)^{\frac{1}{1-\alpha}} .$$

With these preferences, the intertemporal elasticity of substitution for deterministic consumption paths is $1/\alpha$ and the risk aversion to a static gamble is θ. Risk aversion to a dynamic gamble is more complicated, because the intertemporal composition of risk matter: for $\theta > \alpha$, the agent prefers an early resolution of uncertainty. The consumption process in the Barro-Rietz model is:

$$\begin{aligned} \Delta \log C_t &= \mu + \sigma \varepsilon_t, \text{ with probability } 1-p, \\ &= \mu + \sigma \varepsilon_t + \log(1-b), \text{ with probability } p, \end{aligned}$$

where ε_t is *iid* $N(0,1)$. Hence, each period, with probability p, consumption drops by a factor b. The realization of the disaster is *iid* and statistically independent of ε_t at all dates. To allow for recoveries, I consider the following modification of this process: if there was a disaster in the previous period, with probability π, consumption goes back up by an amount $-\log(1-b)$. (Below I allow for more complex dynamics.) Hence, when $\alpha = \theta$ and $\pi = 0$, the model collapses to the Barro-Rietz model; and for $\pi > 0$ there is some possibility of recovery.

When $\pi = 0$, the risk-free rate and equity premium can be found in closed form:

$$\log R_f = \rho + \alpha\mu - (\alpha(\theta-1)+\theta)\frac{\sigma^2}{2} - \log\left(\frac{1-p+p(1-b)^{-\theta}}{\left(1-p+p(1-b)^{1-\theta}\right)^{\frac{\alpha-\theta}{1-\theta}}} \right),$$

$$\log \frac{ER^e}{R^f} = \sigma^2\theta - \log\left(\frac{1-p+p(1-b)^{1-\theta}}{1-p+p(1-b)^{-\theta}}\right).$$

For $p = 0$ or $b = 0$, we obtain the well-known formulas of the lognormal *iid* model, which generate an equity premium puzzle and a risk-free rate puzzle. For $p > 0$ and $b > 0$, the equity premium is increasing in the probability of disasters p and in their size b, and the risk-free rate is decreasing in the probability of disasters or their size (at least if $\alpha \geq \theta$). Because risk-averse agents fear large changes in consumption, a small probability of a large drop of consumption can make the theoretical equity premium large.

For $\pi > 0$, I did not find any useful closed form solutions, but it is easy to solve numerically the model. I will use the same parameter values as Barro, except for the intertemporal elasticity of substitution α, for which I consider a range of possible values. In particular, I use the historical distribution of disasters b instead of a single value.[4] I also follow Barro and assume that government bonds default with probability 0.4 during disasters, and that the recovery rate is $1-b$. With these parameter values, the equity premium is 0.18% without disasters and 5.6% with disasters and no government defaults, and finally 3.5% with disasters and government defaults. Importantly, this result is influenced by the largest historical disasters: if we use exclude from the distribution of b the ten disasters larger than 40% (which all occurred during World War II), the equity premium is reduced to 0.8%.

Figure 2 plots the equity premium as a function of the probability of a recovery, for four different elasticities of substitution: 1/4 (Barro's number), 1/2, 1 and 2. In this computation, the risk aversion θ is kept constant equal to 4. Note that the four lines intersect for $\pi = 0$ since in this case, consumption growth is *iid*, and the IES does not affect the (geometric) equity premium. Perhaps surprisingly, when the IES is low, the equity premium is increased by the possibility of a recovery.

To understand this result, it is useful to recall the present-value identity in the case of $\theta = \alpha$

[4] This modifies slightly the formulas above: an expectation over the disaster size b, conditional on a disaster occuring, must be added to the formulas.

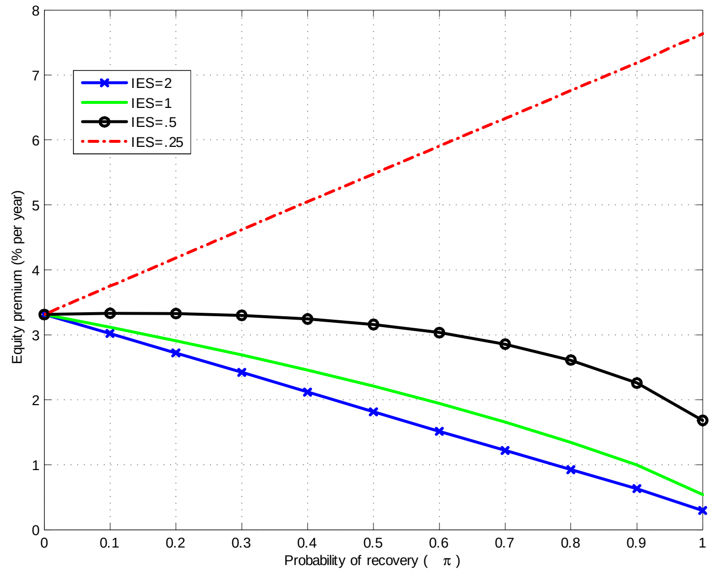

Figure 2: Unconditional, log geometric equity premium, as a function of the IES and the probability of recovery. The risk aversion is $\theta = 4$ and the other parameters are as in Barro (2006).

(state separable utility). The price of a claim to $\{C_t\}$ is

$$\frac{P_t}{C_t} = E_t \sum_{k \geq 1} \beta^k \left(\frac{C_{t+k}}{C_t} \right)^{1-\alpha},$$

hence the fact that a recovery may arise, i.e. that $C_{t+1}, C_{t+2}, ...$, is higher than would have been expected without a recovery, can increase or decrease the stock price today, depending on whether $\alpha > 1$ or $\alpha < 1$. The intuition is that good news about the future have two effects: on the one hand, they increase future dividends (equal to consumption), which increases the stock price today (the cash-flow effect), but on the other hand they increase interest rates, which lowers the stock price today (the discount-rate effect). The later effect is stronger when interest rates rise more for a given change in consumption, i.e. when the intertemporal elasticity of substitution (IES) is low. Given a low IES, the price-dividend ratio falls more when there is a possible recovery than when there are no recoveries. This in turn means that equities are more risky ex-ante, and as a result the equity premium is larger.

Probability of a recovery π	0.00	0.30	0.60	0.90	1.00
IES = 0.25	3.31	4.62	5.91	7.19	7.64
IES = 0.50	3.31	3.30	3.03	2.26	1.68
IES = 1	3.31	2.69	1.94	1.00	0.54
IES = 2	3.31	2.42	1.52	0.63	0.30

Table 2: Unconditional log geometric equity premium, as a function of the intertemporal elasticity of substitution IES 1/alpha and the probability of a recovery pi. This table sets risk aversion theta=4 and the other parameters as in Barro (2006).

When the IES is not low however, recoveries reduce the equity premium. Table 2 summarizes the results. When the IES is equal to one, the equity premium falls by 1/3 if the recovery occurs with 60% probability, and the equity premium is divided by three if the recovery occurs with 90% probability. The intuition is that the decrease in dividends is transitory and thus in disasters stock prices falls by a smaller amount than dividends do, making equities less risky.

This result is consistent with the literature on autocorrelated consumption growth and log-normal processes (John Y. Campbell (1999), Ravi Bansal and Amir Yaron (2004)). While Bansal and Yaron emphasize that the combination of *positively* autocorrelated consumption growth and an IES *above* unity can generate large risk premia, Campbell shows that when consumption growth is *negatively* autocorrelated, risk premia are *larger* when the IES is below unity. Recoveries induce negative serial correlation, so even though Campbell's results are not directly applicable (because the consumption process is not lognormal), the intuition seems to go through.

Clearly the recovery process studied above is too simple: recoveries might not occur right after a disaster, they sometimes also occur more slowly. Moreover, the size of the recovery is uncertain. In Gourio (2007), I consider more general recovery processes. A moderate delay does not affect the results significantly.

Of course, there is no clear agreement on what is the proper value of the IES. The standard view is that it is small (e.g. Robert Hall (1988)), but this has been challenged by several authors (see among others Bansal and Yaron (2004), Fatih Guvenen (2006), Casey Mulligan (2004), Annette Vissing-Jorgensen (2002)). How then, can we decide which IES is more reasonable for the purpose of studying recoveries? The natural solution is to use data on asset prices during disasters. Figures

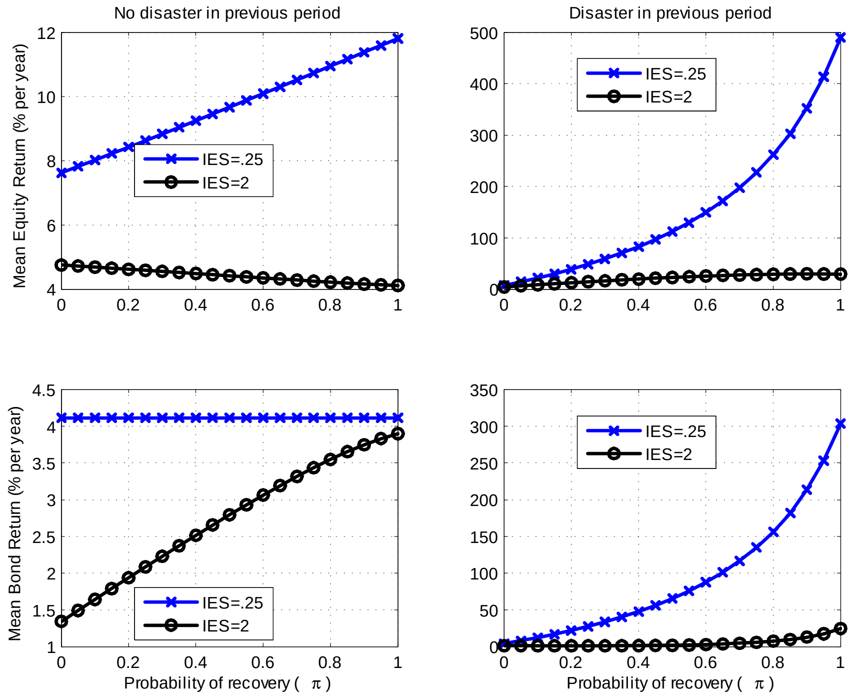

Figure 3: Conditional expected return on equity and bond, as a function of the state, for different values of IES and probability of recovery.

3 and 4 depict the implications of the model with recoveries for two levels of IES (.25 and 2). The baseline disaster model is the case where the probability of recovery is zero. The panels present the expected equity return, risk free rate and risk premium as well as the P-D ratio, conditional on the current state (no disaster in the previous period, or a 35% disaster just occurred).

These figures reveal that when the IES is low (the Barro calibration), a positive probability of recovery implies very large interest rates following a disasters: as consumers are momentarily poor, they want to borrow against their future income, which drives the interest rate up. These huge interest rates are certainly not observed in the data. (Perhaps, people did not anticipate the recoveries.) The high IES case of course implies interest rates which are much smaller, but it also implies that the P-D ratio increases slightly following a disaster (compare the two bottom panels of Figure 4). In contrast, the low IES model implies that the P-D ratio falls. While the P-D probably does not increase following a disaster, it is not necessarily clear that it falls significantly: for instance, according to the Shiller data, the P-E ratio was 20.2 in September 1929 before the

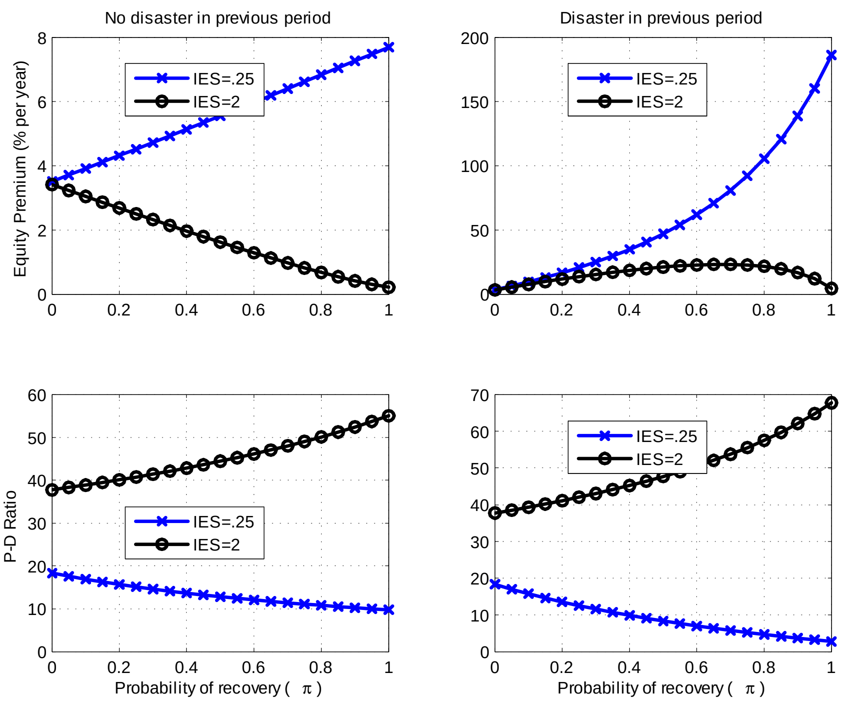

Figure 4: Conditional values of the arithmetic equity premium and P-D ratio, as a function of the state, for various IES and probability of recovery.

crash, 18.4 one year later, and 17.8 in September 1932 at the trough.[5] Obviously earnings fell dramatically, but the question is, Did prices fall more than earnings or dividends? Similarly, dividends fall to zero and even become negative in the Great Depression according to NIPA.

Hence, it is not clear which model fits the data best. If we want to match a low P-D ratio in disasters, an extension of the model seems required - for instance, if people become more fearful in disasters (i.e. increase their estimate of p, perhaps as a result of learning), then the P-D ratio may fall as the equity premium is large, without having a large effect on the interest rate. Many researchers have argued that the Great Depression significantly affected the expectations of households in the following years or decades. A numerical example is the following: assume that, following a disaster, the economy enters a 'waiting' state where the probability of disaster is three times higher than usual. Each period, with probability λ, the economy escapes the waiting state, experiences with probability π a recovery, and return to the normal probability a disaster. I set

[5] These numbers are for current price over current earnings, rather than the smoothed earnings prefered by Shiller. Note also that prices did fall temporarily during the Depression, bu my point is that even at the trough it was not obviously low.

$\alpha = \theta = 4$, $\beta = .99$, and $\lambda = .3$, $\pi = .8$ and for simplicity a single disaster size $b = 0.4$. This economy has the feature that the P-D ratio falls from 24.4 to 18.8 in disasters, the risk-free rate is only 2.8% in disasters (versus 1.6% in normal times), and the equity premium is 4.2% in normal times and 5.8% unconditionally. Hence, introducing a higher likelihood of disasters following a disaster leads to a lower interest rate spike following a disaster.

4 Return Predictability in the Disaster model

Given the relative success of the disaster model in accounting for the risk-free rate and equity premium puzzles, it is important to study if the model can also account for other asset pricing facts, such as the time-series predictability of returns. Empirical research suggests that the excess stock return is forecastable. The basic regression is

$$R_{t+1}^{e} - R_{t+1}^{f} = \alpha + \beta \frac{D_t}{P_t} + \varepsilon_{t+1},$$

where R_{t+1}^{e} is the equity return and R_{t+1}^{f} the risk-free return. As an illustration, John Cochrane (2007) reports for the annual 1926-2004 U.S. sample: $\beta = 3.83$ (t-stat = 2.61, $R^2 = 7.4\%$). A key feature of the data is that using as the left-hand side the equity return R_{t+1}^{e} rather than the excess return $R_{t+1}^{e} - R_{t+1}^{f}$ does not change the results markedly: $\beta = 3.39$ (t-stat = 2.28, $R^2 = 5.8\%$).

To generate variation in expected returns over time, we need to introduce some variation over time in the riskiness of stocks. The natural idea is to make the probability of disaster-time varying. Hence, consider the following environment: there a representative agent who has CRRA utility with risk aversion γ. The disaster probability changes over time according to a monotone first-order Markov process, governed by the transition probabilities $F(p_{t+1}|p_t)$, where p_t is the probability of a disaster at time $t+1$, which is drawn at time t. Formally, $\Delta \log C_{t+1} = \mu + \sigma\varepsilon_{t+1}$ with probability $1 - p_t$, and $\Delta \log C_{t+1} = \mu + \sigma\varepsilon_{t+1} + \log(1 - b)$ with probability p_t. Assume that the realization of p_{t+1} is independent of the realization of disasters at time $t+1$, conditional on

p_t. (This simplification allows to obtain an analytical solution; it implies that the P-D ratio is conditionally uncorrelated with current dividend growth.) The following result is easy to prove:

Proposition 1 *If the probability p_t is always small, then (1) the risk-free rate and expected equity return are decreasing in p_t, (2) the equity premium is increasing in p_t, (3) the P-D ratio is increasing in p_t if and only if $\gamma > 1$.*

This result implies that the correlation between equity risk premia and price-dividend ratio is positive if $\gamma > 1$ and negative (as in the data) if $\gamma < 1$. The intuition is simple: an increase in p reduces expected growth, hence people want to save, which decreases both the risk-free rate and the expected equity return. Because the risk of disaster is higher, the risk premium also increases. The P-D ratio may go up or down, depending on whether the change in expected return is larger than the change in expected dividend growth, which depends on the strength of the interest rate response, and thus on the IES.

This result creates a problem for the simplest model of disasters. First, while most researchers use $\gamma > 1$, this generates a counterintuitive *positive* correlation between the P-D ratio and disaster probability, and this implies that a high P-D ratio forecasts a smaller risk premium, which is the inverse of the data. Hence, we need an IES $1/\gamma$ above unity to generate the key finding of excess return predictability. But using $\gamma < 1$ reduces the risk premium and also implies that recoveries reduce the equity premium.

More fundamentally, there are *no* parameter values which will generate both the stock return *and* excess stock return predictability. The model generates predictability by generating large changes in interest rates, while in the data predictability is due to time-varying risk premia. The natural escape route is to separate the IES and risk aversion, and to use the IES to control movements in the risk-free rate. When the disaster probability is *iid*, i.e. $F(p_{t+1}|p_t) = F(p_{t+1})$, and risk aversion θ is greater than unity, it is possible to show that the result above still applies: the P-D increasing in the probability of a disaster p if and only if the elasticity of substitution is less than one, i.e. $\alpha > 1$. Numerical experiments suggest that relaxing the *iid* assumption for p does not help significantly. As in Bansal and Yaron (2004), it may be necessary to incorporate an additional state variable proxying for time-varying risk.

It is also possible to extend the result above to the case of a time-varying size of disaster b. Formally, assume that $\Delta \log C_{t+1} = \mu + \sigma\varepsilon_{t+1}$, with probability $1-p$, and $\Delta \log C_{t+1} = \mu + \sigma\varepsilon_{t+1} + \log(1-b_t)$, with probability p; that b_t follows a first-order Markov process and that the realization of $\Delta \log C_{t+1}$ and b_{t+1} is independent conditional on b_t. Then:

Proposition 2 *For p low enough, (1) the risk-free rate and expected equity return are decreasing in b_t, (2) the equity premium is increasing in b_t, (3) the P-D ratio is increasing in the size of disaster b_t if and only if $\gamma > 1$.*

Together these two results suggest that the standard calibration of the disaster model with $\gamma > 1$ does not fit the predictability evidence: the model will generate that a high P-D ratio forecasts low expected equity returns; but not that a high P-D ratio forecasts a low excess return on equity. Gabaix (2007) resolves this tension by assuming that the size of dividends disaster changes over time, but not the size of consumption disasters. There may be other resolutions of this conundrum, but they remains to be worked out.

5 Cross-Sectional Implications of the Disaster Model

Empirical research in finance has documented substantial heterogeneity across stocks in expected returns. According to the disaster model, this can be traced to the heterogeneity in responses to 'disasters'. The terrorist attacks of 9-11 offers an interesting example. On 9-17-01 (the first day of trading on the NYSE after 9-11), the S&P 500 dropped 4.9%, but some industries fared very differently: the S&P 500 consumer discretionary index fell 9.8%, the energy index 2.9%, the health care index 0.6%, while defense industry stocks soared: Northrop Grumman was up 15.6% and Lockheed Martin 14.7%.

Formally, starting in the simple Barro-Rietz setup, consider the following dividend process for asset i :

$$\begin{aligned} \Delta \log D_{i,t+1} &= \mu_i + \lambda_i \varepsilon_t \text{ if no disaster,} \\ &= \mu_i + \lambda_i \varepsilon_t + \eta_i \log(1-b) \text{ if disaster.} \end{aligned}$$

Hence, assets differ both in their trend growth μ_i as well as exposure λ_i to 'standard business shocks' and their exposure η_i to disasters. The log risk premium on asset i is:

$$\log\left(\frac{ER_i}{R_f}\right) = \lambda_i \theta \sigma_c + \log\left(\frac{1-p+p(1-b)^{-\theta}}{1-p+p(1-b)^{\eta_i-\theta}}\right),$$

hence the risk premia is determined by the standard exposure to 'business cycle shocks' and a new term, the exposure to disasters η_i. Assets with high η_i will have high average returns. This hypothesis is not easy to test, because it is hard to measure the sensitivity η_i of an asset to disasters. This section presents two preliminary results using the cross-section of stocks: the first one uses the 'natural experiment' of 9-11, and the second one computes the exposure of stocks to large downside market movements. (Clearly, it would be interesting to look at other assets such as bonds or options.)

5.1 The 9-11 Natural Experiment

While 9-11 is not a disaster according to Barro's definition, many people feared at the time that it marked the beginning of a disaster. The heterogeneity across industries in response to 9-11 is impressive and suggests a natural test: if we take these responses to the 9-11 'shock' as proxies for the responses to a true disaster, do they justify the differences in expected returns?

Figure 5 plots the mean monthly excess returns (1970-2004) against the return on 9-17 for the 48 industry portfolios constructed by Fama and French. If the disaster story is true, industries which did well on 9-17 (e.g. defense, tobacco, gold, shipping, coal) should have low average returns,

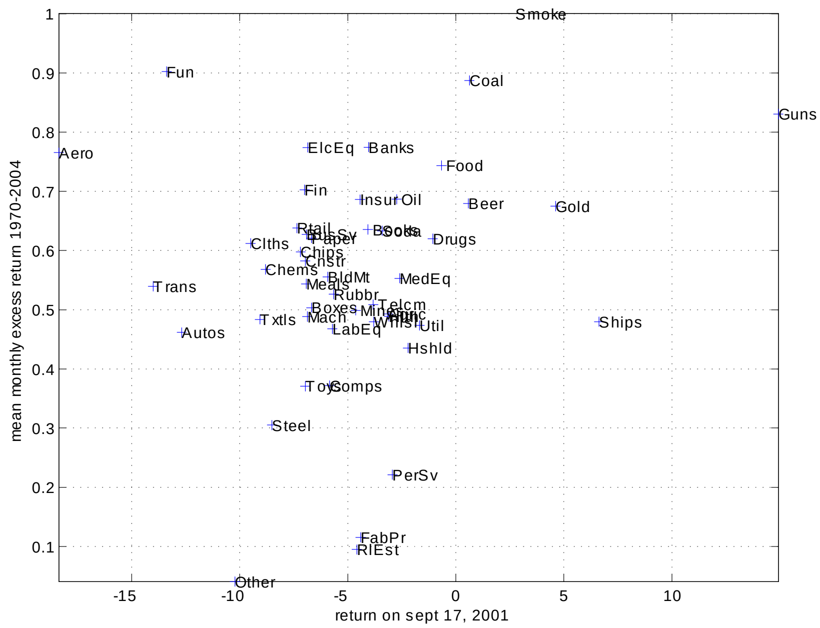

Figure 5: Mean monthly excess return and return on 9-17-01, for 48 portfolios of stocks sorted by industry. Data from prof. French's website.

and industries which did poorly (e.g. transportation, aerospace, cars, leisure) should have high average returns, so we should see a negative relationship. However, the correlation is slightly positive (0.20). Of course one possible answer is that 9-11 is not the ideal experiment, and some industries such as coal or aerospace may have been especially affected by 9-11.

Figure 6 performs the same calculation for the 25 size and book-to-market sorted portfolios. In this case, the correlation is moderately negative (-0.16). Finally the Fama-French factor SMB was up 0.24%, HML was down -0.93%, UMD was up 2.72%, and the small-value/small-growth excess return was -0.20%.[6] Hence, only HML and the small-value/small-growth excess return have the correct sign, and the magnitude is not large.

5.2 Large Downside Risk

The second test is to compute the sensitivity of various portfolios of stocks to large declines in the stock market more generally. Can the disaster explanation account for the cross-sectional puzzles

[6]SMB is a portfolio long small firms and short large firms; HML is long in firms with high book-to-market and short firms in low book-to-market; UMD is long winners (firms with high return in the past month) and short losers. All these strategies generate significant excess returns (see Table 4), which are not accounted for by CAPM or CCAPM betas.

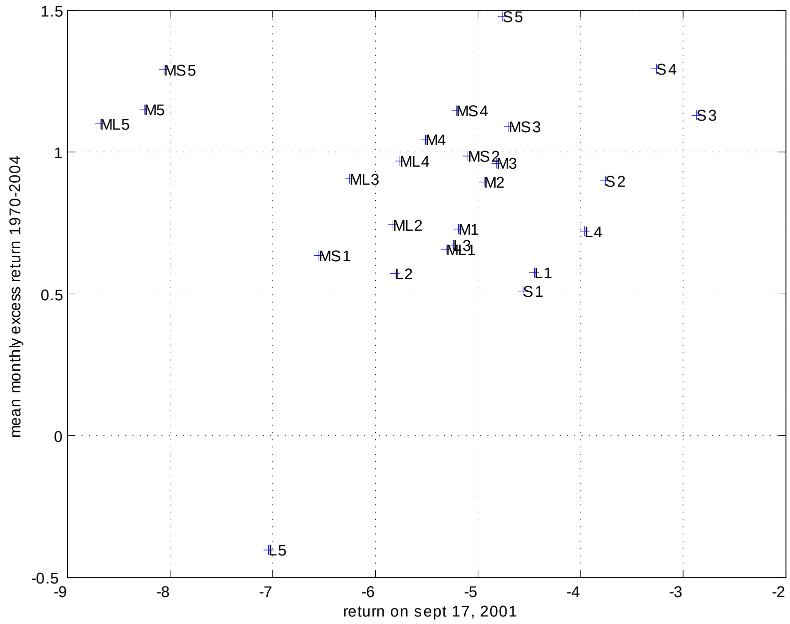

Figure 6: Mean monthly excess return 1970-2004, and return on 9-17-01, for 25 portfolios of firms sorted by size and book-to-market. Data from prof. French's website.

like value-growth, momentum, small-big which have attracted so much attention in the empirical finance literature?

I measure the exposure of assets to large negative events by running the following time series regression:

$$R^i_{t+1} - R^f_{t+1} = \alpha_i + \beta^d_i \left(R^m_{t+1} - R^f_{t+1} \right) \times 1_{\left(R^i_{t+1} - R^f_{t+1}\right) < t} + \varepsilon_{it+1},$$

where R^m_{t+1} is the market return and R^f_{t+1} is the risk-free return. The only change between this model and the CAPM is that the risk factor is the stock market return conditional on a large negative return. Securities which have a large β^d_i do badly when the stock market does very badly. (This can be justified as an approximation to the model above, since the market return in this case is proportional to consumption growth.) I use monthly data and set arbitrarily $t = -10\%$.[7]

When I consider the 25 Fama-French portfolios sorted by size and book-to-market (sample: 1932-2005, with 22 'disaster months'), I find that this model does not improve on the basic CAPM, because the 'disaster beta' has a correlation over .96 with the standard market beta. The same is

[7] Since 1926 there have been 29 months where the excess market return is less than -10%.

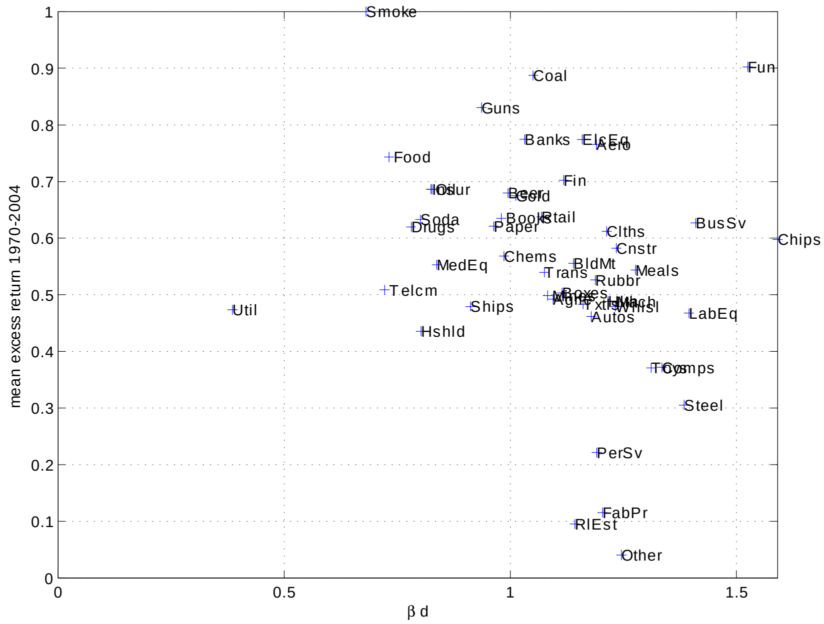

Figure 7: Mean excess return (1970-2004) against the disaster beta, computed from the time series regression (4), for the 48 industry-sorted portfolios.

true when I use the 48 industry portfolios (sample: 1970-2005, with 10 'disaster months'); in this case the correlation between the disaster beta and the standard beta is 0.90, and figure 7 shows clearly that the relation between disaster beta and average return does not exist in these data.

In Table 3, I perform the same regressions for HML, SMB, UMD and small-value-small-growth. We see that the coefficient β^d does not explain very well the mean returns: for UMD, it has the wrong sign, it is insignificant for SV-SG; for HML it is small and borderline significant. Only for SMB is there some empirical support for the disaster story: small firms have indeed more negative returns than large firms when there is a big negative stock return.

It is plainly not clear that value stocks do worse in large negative events. Figure 8 shows that there is little discernible difference between the small-growth and small-value portfolios during the Great Depression. Hence, while small value stocks have higher average returns, it does not appear that they were much more sensitive to extreme events, as measured by the 1929-1932 crash. Another simple way to measure the exposure to disasters is to look at the average return during all the months since 1926 when the monthly excess return on the stock market was less than

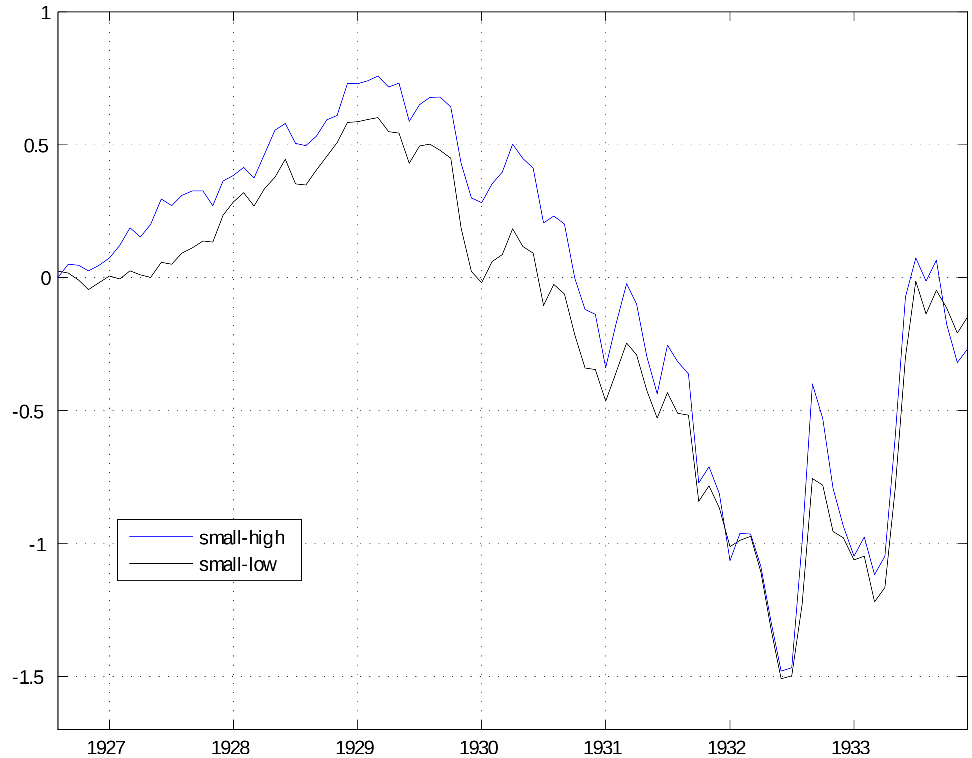

Figure 8: Cumulated geometric return during the Great Depression, for the small value and small growth portfolios of Fama and French. This is $\sum_{k=0}^{t} \log(1 + r_k)$.

-15%; there are 9 such months (seven from 1929 to 1940, October 1987 and August 1998). These moments are reported in Table 4. In five times out of nine, the small-value portfolios outperformed the small-growth portfolios; in five times out of nine, the HML return was positive; in eight out of nine, the momentum excess return UMD is positive. There is some supportive evidence for the SMB asset (small stocks minus large stocks) which return was negative in seven out of nine events.

The main problem with this section is that measuring the exposure to disasters is hard. However, the preliminary conclusion is that there is little support in the data for disaster explanation when looking at the cross-section of stock returns (value, momentum, industries), with the exception of size effects.

	β^d	t-stat
HML	0.08	1.9
SMB	0.18	4.6
UMD	-0.24	4.3
SV-SG	0.02	0.4

Table 3: Beta on Large Negative Returns, for four excess returns.

	$E(R)$	$E(R\|R^m < -.1)$	$E(R\|R^m < -.15)$
HML	0.40	-0.68	0.15
t-stat	3.47	-0.53	0.09
SMB	0.24	-2.69	-2.68
t-stat	2.19	-3.63	-1.90
UMD	0.76	4.26	5.97
t-stat	5.01	3.24	2.56
SV-SG	0.49	0.48	0.46
t-stat	4.14	0.41	0.33

Table 4: Mean returns on the HML, SMB, UMD and small growth-small value portfolios, for the full sample, the sample of market declines greater than ten percent.

6 Conclusion

The disaster explanation of asset prices is attractive on several grounds: first, there are 'reasonable' calibrations which can generate a sizeable equity premium. Second, disasters can easily be embedded in standard macroeconomic models. Moreover, the explanation is consistent with the empirical finance literature which documents deviations from log-normality ('fat tails'). Inference about extreme events is hard, so it is possible that investors' expectations do not equal an objective probability, due to learning and/or a concern for misspecification.

But precisely because the disaster explanation is not rejected on a first pass, we should be more demanding, and study if it can account quantitatively for other asset pricing puzzles, and whether it is robust to reasonable extensions such as recoveries. The current paper points toward some areas which would benefit from further study.

References

[1] Bansal Ravi and Amir Yaron, 2004: "Risks for the Long Run: A Potential Explanation of Asset Pricing Puzzles", Journal of Finance, 59: 1481-1509.

[2] Barro Robert, 2006: "Rare Disasters and Asset Markets in the Twentieth Century", Quarterly Journal of Economics 121:823-866.

[3] Campbell, John, 1999: "Asset Prices, Consumption, and the Business Cycle", Chapter 19 in John B. Taylor and Michael Woodford eds., Handbook of Macroeconomics, Volume 1, North-Holland: Amsterdam, 1231-1303.

[4] Cochane John, 2007: "The Dog that did not Bark", Review of Financial Studies, forthcoming.

[5] Epstein, Larry and Stanley Zin, 1989: "Substitution, Risk Aversion, and the temporal Behavior of consumption and asset returns: A theoretical framework", Econometrica 57: 937-969.

[6] Gourio, François, 2007. "Disasters, Recoveries, and Predictability". Mimeo, Boston University.

[7] Guvenen, Fathi, 2006: "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective", Journal of Monetary Economics 53: 1451-1472.

[8] Hall, Robert, 1988: "Intertemporal Substitution in Consumption", Journal of Political Economy 96: 339-357.

[9] Hansen, Lars and Kenneth Singleton, 1982: "Generalized instrumental variables estimation of nonlinear rational expectations models", Econometrica 50:1269-1286.

[10] Liu, Jun, Jun Pan and Tan Wang, 2005: "An Equilibrium Model of Rare-Event Premia and Its Implication for Option Smirks", Review of Financial Studies: 18:131–164.

[11] Maddison Angus, 2003: The World Economy: Historical Statistics, Paris: OECD, 2003.

[12] Mulligan Casey, 2004: "Capital, Interest and Aggregate Intertemporal Substitution", Working Paper, University of Chicago.

[13] Ranciere, Romain, Aaron Tornell and Frank Westermann 2007: "Systemic Crises and Growth", Quarterly Journal of Economics, forthcoming.

[14] Rietz Thomas, 1988: "The Equity Premium: a Solution", Journal of Monetary Economics 22: 117-131.

[15] Rodriguez, Carlos, 2006: "Consumption, the persistence of shocks, and asset price volatility", Journal of Monetary Economics, 53(8): 1741-1760.

[16] Vissing-Jorgensen Annette, 2002: "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution" 2002.", Journal of Political Economy 110: 825:853.