

Disaster risk and its implications for asset pricing*

Jerry Tsai

Jessica A. Wachter

University of Oxford

University of Pennsylvania

and NBER

January 26, 2015

Abstract

After laying dormant for more than two decades, the rare disaster framework has emerged as a leading contender to explain facts about the aggregate market, interest rates, and financial derivatives. In this paper we survey recent models of disaster risk that provide explanations for the equity premium puzzle, the volatility puzzle, return predictability and other features of the aggregate stock market. We show how these models can also explain violations of the expectations hypothesis in bond pricing, and the implied volatility skew in option pricing. We review both modeling techniques and results and consider both endowment and production economies. We show that these models provide a parsimonious and unifying framework for understanding puzzles in asset pricing.

*Tsai: Department of Economics, Oxford University and Oxford-Man Institute of Quantitative Finance ; Email: jerry.tsai@economics.ox.ac.uk; Wachter: Department of Finance, The Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104; Tel: (215) 898-7634; Email: jwachter@wharton.upenn.edu. We thank Robert Barro, Xavier Gabaix, and Mete Kilic for helpful comments.

1 Introduction

Many asset pricing puzzles arise from an apparent disconnect between returns on financial assets and economic fundamentals. A classic example is the equity premium puzzle, namely that the expected return on stocks over government bills is far too high to be explained by the observed risk in consumption. Another famous example is the volatility puzzle. Economic models trace stock return volatility to news about future cash flows or discount rates. Yet dividend and interest rate data themselves provide very little basis for explaining this volatility. Other puzzles can be stated in similar ways.

In this paper, we survey a class of explanations for these and other findings known as rare disaster models. These models assume that there is a small probability of a large drop in consumption. Under the assumption of isoelastic utility, this large drop is extremely painful for investors and has a large impact on risk premia, and, through risk premia, on prices. The goal of these models, broadly speaking, is to take this central insight and incorporate it into quantitative models with a wide range of testable predictions.

How does the risk of rare disasters differ from the risk that has already been carefully explored in models of asset pricing over the last quarter century? What differentiates a rare disaster from merely a large shock? Crucial to understanding of rare disasters is to understand when they have *not* occurred, namely in the postwar period in the United States. A central contention of the rare disaster literature is that the last 65 years of U.S. data has been a period of calm that does not represent the full spectrum of events that investors incorporate into prices. This already says something about the probabilities of rare disasters: for example, if the probability were 10% a year, then the probability of having observed no disasters is 0.90^{65} , or 0.001. While we may be very lucky in the United States, and while there may be survival bias in the fact that most finance scholars find themselves situated in the U.S. and studying U.S. financial markets, assuming this much luck and thus this much survival bias would be unsettling. Thus historical time series says that disasters, if they exist, must be rare.

A second aspect of rare disasters is that they must be large, so that their size essentially rules out the normal distribution, just as the discussion above essentially rules out disasters that occur with probability 10%. Thus the assumption of rare disasters is akin to rejecting the assumption of normality, while still allowing for the fact that data might plausibly appear to be normal over a time span that lasts as long as 65 years.

This rejection of the normal distribution sounds like a technical condition, divorced from underlying economics. However, it is of fundamental importance to understanding why one obtains qualitatively different results from incorporating rare disasters into economic models. Under the normal distribution, risk is easy to measure because the normal distribution implies that small and thus frequently observed changes in the process of interest can be used to make inferences about large and thus infrequently observed changes in prices and fundamentals. Stated in econometric terms, normality is a strong identifying assumption. If one seeks to measure risk, and the object of interest has the normal distribution, then a small amount of data (65 years is more than enough) can be used to obtain a very accurate estimate of the risk. The assumption of normality, while not universal, is ubiquitous in finance. For example, the diffusion process is the standard model used in continuous-time finance. Yet a diffusion is nothing more than a normal distribution with parameters that may themselves vary according to normal distributions. And while the assumption of the normal distribution is often clearly stated, its consequence, that small changes in a process are informative of larger changes, is not.

In contrast to the situation in standard models, risk in rare disaster models is difficult to measure. Small variations in a process of interest are not informative of large variations. Thus investors may rationally believe a disaster is within the realm of possibility, even after many years of data to the contrary. This insight leads to a rethinking of the connections between stock prices and fundamentals, as shown by the models that we survey.

The rest of this survey proceeds as follows. Section 2 presents a general result on risk premia in models of disaster risk. This result is a building block for many other results that follow. Section 3 presents the basic model of consumption disasters that is used to explain

the equity premium puzzle, and recent extensions. Section 4 discusses the dynamic models that can explain the volatility puzzle. Section 5 discusses extensions to fixed income and to options. Finally, Section 6 reviews the literature on disaster risk in production-based models. These models endogenize consumption and cash flows and offer a deeper understanding of the links between stock prices and the real economy.

2 Disaster Risk Premia

Most asset pricing puzzles, in one way or another, come down to a difference between a modeled and observed risk premium (the expected return on a risky asset above a riskless asset). This is obvious when the puzzle refers to the mean of a return. When the puzzle refers to a volatility or covariance it is also true, since time-variation in risk premia appear to be a major determinant of stock price variation. For this reason, we begin our survey with a theorem on risk premia in models of disaster risk.¹

Consider a model driven by normal risk, modeled by a vector of Brownian (or diffusive) shocks B_t , and by rare event risk, modeled by a Poisson shock N_t . The Poisson shock has intensity λ_t , meaning that it is equal to 1 with probability λ_t and zero otherwise.² This basic jump-diffusion model was introduced by Merton (1976) for the purpose of calculating option prices when stock prices are discontinuous.

We assume a process for state prices. This process can either be interpreted as the marginal utility of the representative agent (as will be the case in the models that follow), or as simply the process that is guaranteed to exist as long as there is no arbitrage (Harrison and Kreps (1979)). The state-price density is given by π_t and follows the process

$$\frac{d\pi_t}{\pi_{t-}} = \mu_\pi dt + \sigma_\pi dB_t + (e^{Z_\pi} - 1) dN_t. \quad (1)$$

¹This result unifies results already reported in the literature, for, e.g. iid models of stock prices (Barro (2006)), dynamic models of stock prices (Longstaff and Piazzesi (2004) and Wachter (2013), and exchange rate models (Farhi and Gabaix (2014)).

²More precisely, the probability of k jumps over the course of a period τ is equal to $e^{-\lambda\tau} \frac{(\lambda\tau)^k}{k!}$. We take τ to be in units of years.

The parameters μ_π and σ_π could be stochastic (here and in what follows we assume B_t is a row vector so whatever multiplies it, in this case σ_π , is a column vector). The ratio of π_t to itself at different points of time is commonly referred to as the stochastic discount factor. The change in $\log \pi_t$, should a rare event occur, is Z_π , a random variable which we assume for now has a time-invariant distribution. The reason to model the change in $\log \pi_t$ rather than π_t , is to ensure that prices remain positive. Disasters are times of low consumption, high marginal utility, and therefore high state prices. Thus we emphasize the case of $Z_\pi > 0$.

The asset (we refer to this as a “stock” to be concrete, but it need not be) with price process

$$\frac{dS_t}{S_{t-}} = \mu_S dt + \sigma_S dB_t + (e^{Z_S} - 1) dN_t. \quad (2)$$

Again, μ_S and σ_S can be stochastic. The change in the log stock prices in the event of a disaster is given by the random variable Z_S , which can be correlated with Z_π and has a time-invariant distribution. Once again, the model is written so that the price cannot go negative. Most cases of interest involve a price that falls during a disaster, which corresponds to $Z_S < 0$. We will use the notation E_ν to denote expectations taken with respect to the joint distribution of Z_S and Z_π .

Let D_t denote the continuous dividend stream paid by the asset (this may be zero). Then the expected return (precisely, the continuous-time limit of the expected return over a finite interval) is defined as

$$r_t^S = \mu_S + \lambda_t E_\nu [e^{Z_S} - 1] + \frac{D_t}{S_t},$$

which includes the drift in the price, the expected change in price due to a jump, and the dividend yield. Proposition 1 gives a general expression for risk premia under these assumptions.

Proposition 1. *Assume no-arbitrage, with state prices given by (1). Consider an asset specified with price process (2). Let r_t denote the continuously compounded riskfree rate.*

Then continuous-time limit of the risk premium for this asset is

$$r_t^S - r_t = -\sigma_\pi \sigma_S^\top - \lambda_t E_\nu [(e^{Z_\pi} - 1)(e^{Z_S} - 1)]. \quad (3)$$

The expected return over the riskfree rate in periods without disasters will be

$$r_t^S - r_t + \lambda_t E_\nu [e^{Z_S} - 1] = -\sigma_\pi \sigma_S^\top - \lambda_t E_\nu [e^{Z_\pi} (e^{Z_S} - 1)]. \quad (4)$$

Consider the first equation in Proposition 1, the risk premium on the asset with price S_t . There are two terms: the first corresponds to the Brownian risk, and it takes a standard form (see Duffie (2001, Chapter 6)). The second corresponds to Poisson risk. Because the Brownian shocks and Poisson shocks are independent, we can consider the two sources separately. Both terms represent covariances. The second term is the covariance of prices and marginal utility in the event of a disaster, multiplied by the probability that a disaster occurs.³ For an asset that falls in price when a disaster occurs, the product $(e^{Z_\pi} - 1)(e^{Z_S} - 1)$ is negative. Such an asset will have a higher expected return because of its exposure to disasters.

Now consider the second equation in Proposition 1. This is the expected return that would be observed in samples in which disasters do not occur. It is equal to (3), minus the expected price change in the event of a disaster. Because the price falls in a disaster, (4) will be greater than the true risk premium: observing samples without disasters leads to a bias, as pointed out by Brown, Goetzmann, and Ross (1995) and Goetzmann and Jorion (1999).

In what follows, we will quantitatively evaluate the sizes of the disaster term in the true risk premium (3), and its relation to the observed risk premium (4) in economies where stock prices are determined endogenously. We will also show that disasters have a role to play in the Brownian terms σ_π and σ_S . If, for example, the probability of a disaster varies over time, then this will in general be reflected in σ_π and σ_S . Thus the equations above, taken at

³Because the expected change in both quantities over an infinitesimal interval is zero, the expected co-movement between these quantities is in fact the covariance.

face value, may understate the role of disasters.

3 Consumption Disasters

Two central puzzles in asset pricing are the equity premium puzzle (Mehra and Prescott (1985)) and the riskfree rate puzzle (Weil (1989)). Barro (2006) shows that, when disasters are calibrated using international GDP data, disaster risk can resolve both puzzles. Here, we consider a continuous-time equivalent of Barro’s model, in which we make use of a now-standard extension of time-additive isoelastic utility that separates risk aversion from the inverse of the elasticity of intertemporal substitution (Epstein and Zin (1989), Duffie and Epstein (1992)).⁴ This extension, because it is based on isoelastic utility, implies that risk premia and riskfree rates are stationary even as consumption and wealth grow, just as in the actual economy. This scale-invariance will make calibration possible. Details of the utility specification are contained in Appendix B.

In what follows, the parameter β refers to the rate of time preference, γ to relative risk aversion, and ψ (in an approximate sense) to the elasticity of intertemporal substitution (EIS).⁵

3.1 The model

Our model for consumption follows along the same lines as the model considered in Section 2:

$$\frac{dC_t}{C_{t-}} = \mu dt + \sigma dB_t + (e^{Z_t} - 1) dN_t. \quad (5)$$

As in Section 2, B_t is a standard Brownian motion and N_t is a Poisson process with constant intensity λ . Z_t is a random variable whose time-invariant distribution ν is independent of N_t and B_t . A disaster is represented by $dN_t = 1$, and e^{Z_t} is the change in consumption

⁴See Martin (2013) for a solution to a general class of models of which this model is an example.

⁵We follow standard terminology in referring to ψ as the EIS. However, in the models we consider, uncertainty implies that there is not a simple mapping between ψ and the sensitivity of consumption to interest rates.

should a disaster occur. Thus disasters are represented by $Z_t < 0$; using an exponential ensures that consumption stays positive. In this model consumption growth is independent and identically distributed over time. One advantage of this model is that it allows for long periods of low volatility and thin tails (like the post-war period in the United States), along with large consumption disasters (as represented by the Great Depression).

The model is not complete without a specification for the dividend process for the aggregate market. The simplest assumption is that dividends equal consumption. However, it is sometimes convenient to allow the aggregate market to differ from the claim to total wealth. We will follow Abel (1999) and Campbell (2003) in assuming

$$D_t = C_t^\phi, \tag{6}$$

though for much of the discussion in this section, ϕ will be restricted to 1. The parameter ϕ is sometimes referred to as “leverage”. The specification (6) allows us to take into account the empirical finding that dividends are procyclical in normal times and fall far more than consumption in the event of a disaster (Longstaff and Piazzesi (2004)). Ito’s Lemma for jump-diffusion processes (Duffie (2001, Appendix F)) then implies

$$\frac{dD_t}{D_{t-}} = \mu_D dt + \phi\sigma dB_t + (e^{\phi Z_t} - 1) dN_t, \tag{7}$$

where $\mu_D = \phi\mu + \frac{1}{2}\phi(\phi - 1)\sigma^2$.

Let S_t denote the price of the dividend claim. In this iid model, the price-dividend ratio S_t/D_t is a constant. Note that a constant price-dividend ratio implies that the percent decline in prices is equal to the percent decline in dividends, so the process S_t is given by:

$$\frac{dS_t}{S_{t-}} = \mu_S dt + \phi\sigma dB_t + (e^{\phi Z_t} - 1) dN_t \tag{8}$$

for a constant μ_S which we leave unspecified for now.

In this iid model with isoelastic utility, the marginal utility process π_t is proportional to

$C_t^{-\gamma}$, where γ represents relative risk aversion.⁶ Thus, by Ito's Lemma,

$$\frac{d\pi_t}{\pi_{t-}} = \mu_\pi dt - \gamma\sigma dB_t + (e^{-\gamma Z_t} - 1) dN_t. \quad (9)$$

Again, we leave the constant μ_π unspecified for the moment. Thus marginal utility jumps up if a disaster occurs. The greater is risk aversion, the greater the jump.

3.2 The equity premium

Using Proposition 1 we can derive the equity premium in the model:

$$r^S - r = \gamma\phi\sigma^2 - \lambda E_\nu [(e^{-\gamma Z_t} - 1)(e^{\phi Z_t} - 1)]. \quad (10)$$

Equation (10) contains two terms. The first, $\gamma\phi\sigma^2$ is the basic consumption CAPM term (Breenden (1979)). The second, as explained in Section 2, is due to the presence of disasters, and is positive because a disaster is characterized by falling consumption (and thus rising marginal utility) and falling prices.

For the remainder of this section, we will follow Barro (2006) and Rietz (1988) and specialize to the unlevered case of $\phi = 1$. Thus our target for the equity premium is the “unlevered” premium, which can reasonably be taken to be 4.8% a year (Nakamura, Steinsson, Barro, and Ursúa (2013)). Note that thus far, only risk aversion appears in (10), and this equation is identical to what we would find in the time-additive utility case.

In writing down the process for consumption (5), we have split risk into two parts, a Brownian part and a rare disaster part. The equation for the risk premium (10) reflects this separation. Consider two ways of interpreting U.S. consumption data from 1929 to the present.⁷ First assume that consumption growth is lognormally distributed, namely (5), with the jump component set to zero. We can then compute σ^2 based on the measured volatility of this series, which is about 2.16% per annum. With an unlevered equity premium of 4.8%,

⁶For convenience, we provide a proof of this standard result in the Online Appendix.

⁷The numbers in this section refer to real per capita annual consumption growth, available from the Bureau of Economic Analysis.

γ would have to be about 102 to reconcile the equity premium with the data. This is the equity premium puzzle.

Now consider a different way to view the same data. In this view, the Great Depression represents a rare disaster. Consumption declined about 25% in the Great Depression. Assuming a 1% probability of such a disaster occurring (and, for the moment, disregarding the tiny first term), explaining an equity premium of 4.8% only requires $\gamma = 10.5$. This is the point made by Rietz (1988) in response to Mehra and Prescott (1985). Many would say that a risk aversion of 10.5 is still too high, and in fact more recent calibrations justify the equity premium with a lower risk aversion coefficient, but without a doubt, this simple reconsideration of the risk in consumption data has made a considerable dent in the equity premium.

Which way of viewing the consumption data is more plausible? Given the low volatility of consumption over the 65-year postwar period (1.25%), a constant volatility of 2.16% is virtually impossible, as can be easily seen by running Monte Carlo simulations of 65 years with normally distributed shocks. If consumption were normally distributed, observing a sample 65 years long with a volatility of 1.25% occurs with probability less than one in a million.

What about the rare disaster model? Here, Mehra and Prescott (1988), and, more recently, Constantinides (2008) and Julliard and Ghosh (2012) have argued that this too is improbable. The problem is that consumption did not decline by 25% instantaneously in the Great Depression, but rather there were several years of smaller declines. Our equations, however, assume that the declines are instantaneous. Does this matter? An interesting difference between jump, or Poisson, risk and Brownian, or normal risk, is that it might matter. Normally distributed risk (assuming independent and identically distributed shocks) is the same regardless of whether one measures it over units of seconds or of years. Not so in the case of Poisson risk. As the Great Depression took 4 years to unfold, suppose instead of a 1% chance of 25% decline, we had a 4% chance of a 7% decline.⁸ In this case, the population

⁸A total decline of 25% implies $Z = -0.288$. For this exercise, we divide Z by 4, to find $Z = -0.072$. This implies a percentage decline of 6.9%.

equity premium would only be 0.31%, while the observed equity premium would be 0.59%; in other words an order of magnitude lower. These authors argue that an all-at-once 25% decline, given U.S. consumption history, is too unlikely to be worth considering.

As we will see, however, it matters very much whether one models the consecutive declines as iid jumps, and whether one requires time-additive utility. In a sense, time-additive utility is a knife-edge case, in which the only source of risk that matters in equilibrium is instantaneous shocks to consumption growth. When the agent has a preference over early resolution of uncertainty, and shocks are not iid but rather cluster together, the model once again produces equity premia of similar magnitude as if the decline happened all at once. We will show this in a simple model in Section 3.7.

Clearly, calibration of the disaster probability λ_t and the disaster distribution Z_t is key to evaluating whether the model can explain the equity premium puzzle. We now turn to calibration of these parameters.

3.3 Calibration to international macroeconomic consumption data

Following the work of Rietz (1988), rare disasters as an explanation of the equity premium received little attention for a number of years (exceptions include Veronesi (2004) and Longstaff and Piazzesi (2004)). This changed with the work of Barro (2006), who calibrates the model above to data on GDP declines for 35 countries over the last century (the original data are from Maddison (2003)). A disaster is defined as a cumulative decline in GDP of over 15%. Altogether he finds 60 occurrences of cumulative declines of over 15%, for the 35 countries over the 100 years. This implies a probability of a disaster of 1.7%. These declines also give a distribution of disaster events. This size distribution puts the Great Depression in a different light. The average value is 29%, about the same as the 25% decline assumed above. However, the presence of much larger declines in the sample, as large as 64%, combined with isoelastic utility, implies that the model behaves closer to the case of a 50% decline.

Barro (2006) uses GDP data; consumption data are closer to aggregate consumption in the model. In normal times, GDP is considerably more volatile than consumption, so

calibrating a consumption-based economy to GDP data makes puzzles (artificially) easier to solve. Barro and Ursúa (2008) build a dataset consisting of consumption declines across these, and additional countries. They also broaden and correct errors in the GDP dataset of Maddison (2003). They find that it matters little for the equity premium whether consumption or GDP is used. This is because the largest disasters, which drive the equity premium result, occur similarly in consumption and GDP. In fact, in wars, GDP tends to decline by less because it is buffered by war-time government spending.

Figure 1 shows the histogram of consumption declines across the full set of countries (Panel A) and OECD countries (Panel B). There is significant mass to the right of 25%; it is these larger disasters that make it possible to account for the equity premium with risk aversion as low as 4. The large numbers to the right represent the effects of World War II in Europe and Asia.⁹ The rare disaster model embeds into U.S. equity prices a positive probability (albeit very small) of a consumption decline of this magnitude.

The disasters observed in these data are not independent of course: Most of the events, particularly in OECD countries can be associated with either World War I, World War II, or the Great Depression. Does this matter? If we were to formally estimate the probability of a binomial model of a disaster/no disaster state, it might matter for the standard errors. However, a 2% probability does not seem unreasonable even if one took the view that there were three world wide disasters in the course of 100 years (or perhaps two...). Does the lack of independence matter for the interpretation of the histogram in Figure 1? We could take the view that we should look at a wealth-weighted consumption decline across all the countries for the three major disasters. Leaving aside the difficulty of constructing such a measure, we are then still left with essentially three observations, which does not seem enough to rule out the histogram in Figure 1.¹⁰ One view of these data is they provide a disciplining device for

⁹This histogram is not exhaustive of the misfortunes that have befallen mankind relatively recently that might be priced into equity returns. For example, the data include neither the experience of the American South following the Civil War, nor the experience of Russia throughout this century.

¹⁰A related question is whether it is appropriate to apply these international data to the United States. Given our assumed parameters, the U.S. experience is far from an anomaly, as Nakamura, Steinsson, Barro, and Ursúa (2013) discuss. These data can neither prove nor disprove that the U.S. is subject to this distribution in the minds of investors; this is a matter for prior beliefs.

what would otherwise constitute a parameter for which we as a profession have very little information. Nakamura, Steinsson, Barro, and Ursúa (2013) estimate and solve a model in which independence of disasters, duration and partial recovery from disasters are taken into account; they find that the model can still explain the equity premium puzzle for moderate values of risk aversion.

3.4 Interest rates

We now return to the model of Section 3.1 and consider the interest rate. The presence of disasters also affects agent's desire to save. Let β denote the rate of time preference and ψ the elasticity of intertemporal substitution (see Appendix B). The instantaneous riskfree rate implied by investor preferences and the consumption process (5) is

$$r = \beta + \frac{1}{\psi}\mu - \frac{1}{2}\left(\gamma + \frac{\gamma}{\psi}\right)\sigma^2 + \lambda E_{\nu}\left[\left(1 - \frac{1}{\theta}\right)(e^{(1-\gamma)Z_t} - 1) - (e^{-\gamma Z_t} - 1)\right]. \quad (11)$$

This riskfree rate is perhaps comparable to the real return on Treasury Bills (more on this below), so the left hand side is about 1.2% in the data. From the point of view of traditional time-additive utility ($\gamma = 1/\psi$ and $\theta = 1$) with no disasters, this is a puzzle. Postwar consumption growth is about 2% a year; combined with a reasonable level of risk aversion, say, 4, the second term in (11) is about 8%. The third term is negligible for reasonable parameter values because σ^2 is extremely small. Thus β , the rate of time preference, would have to be significantly negative to explain the Treasury Bill rate in the data (Weil (1989)). Unlike with the equity premium puzzle, the separation between the EIS and risk aversion can be helpful here, though it is questionable whether, in this model at least, it can completely resolve the puzzle (Campbell (2003)).

The term multiplying λ in (11) represents the effect of disasters. Even a small probability of a disaster can lead to a much lower riskfree rate because disaster states are very painful for agents, and, in this model at least, savings in the riskfree asset offers complete insurance.¹¹

¹¹Disasters always decrease the riskfree rate. The term $1 - \frac{1}{\theta}$ is bounded above by $\gamma/(\gamma - 1)$, and properties of the exponential imply that $\frac{1}{\gamma}(e^{-\gamma Z_t} - 1) > \frac{1}{\gamma - 1}(e^{-\gamma Z_t} - 1)$.

Table 1 reports riskfree rates implied by risk aversion of 4, for various values of the elasticity of intertemporal substitution. Investors who are more willing to substitute over time are less sensitive to the presence of disasters, all else equal. However, for all values we consider, the presence of disasters implies that the riskfree rate is actually negative.

This discussion implies that investors can completely insure against disasters by purchasing Treasury Bills. Historically, some consumption disasters have coincided with default on government debt, either outright or through inflation. The fact that this was not the case in the Great Depression, does not mean that investors are ruling it out. Following Barro (2006), we can account for default by introducing a shock to government debt in the event of a disaster.¹²

To introduce default, let \mathcal{L}_t denote the price process resulting from rolling over short-term government claims. The only risk associated with this claim is default in the event of a disaster, so

$$\frac{d\mathcal{L}_t}{\mathcal{L}_t} = r^{\mathcal{L}} dt + (e^{Z_{\mathcal{L},t}} - 1) dN_t, \quad (12)$$

where $r^{\mathcal{L}}$ is the return on government bills if there were no default. We capture the relation between default and consumption declines in disasters as follows:

$$Z_{\mathcal{L},t} = \begin{cases} Z_t & \text{with probability } q \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

This implies that in the event of a disaster, there is a probability q of default, and if this happens, the percent loss is equal to the percent decline in consumption. We can then apply Proposition 1 to write down the instantaneous risk premium on this security:

$$r^b - r = -\lambda q E_{\nu} [(e^{-\gamma Z_t} - 1) (e^{Z_t} - 1)]. \quad (14)$$

¹²In this case, the assumption of complete markets still implies that there exists a riskfree rate, it just isn't comparable to the government bill rate. What happens if markets are incomplete and there is no riskfree rate? This is a hard question to answer, because the representative investor framework no longer applies, and we are not aware of any work that addresses it. Intuition suggests that this would make the required compensation for disasters higher rather than lower.

Finally, by definition we have $r^b = r^{\mathcal{L}} + \lambda q E_{\nu}[e^{Z_t} - 1]$, and so the observed premium on government debt in samples without disasters is given by

$$r^{\mathcal{L}} - r = -\lambda q E_{\nu} [e^{-\gamma Z_t} (e^{Z_t} - 1)]. \quad (15)$$

If we restrict to the case with no leverage, comparing (14) with (10) implies an equity premium of

$$r^S - r^b = \gamma \sigma^2 - (1 - q) \lambda E_{\nu} [(e^{-\gamma Z_t} - 1)(e^{Z_t} - 1)]. \quad (16)$$

Furthermore, in samples without disasters, the observed observed equity premium equals

$$\text{observed } r^S - r^{\mathcal{L}} = \gamma \sigma^2 - (1 - q) \lambda E_{\nu} [e^{-\gamma Z_t} (e^{Z_t} - 1)]. \quad (17)$$

The next section combines the results so far to discuss the implications for interest rates and the equity premium in a calibrated economy.

3.5 Results in a calibrated economy

Table 1 puts numbers to the equity premium and government bill rate using international consumption data of Barro and Ursúa (2008). As already discussed, the presence of disasters has dramatic effects on the level of the riskfree rate. Introducing default actually helps the model explain the level of the government bill rate in the data.

Table 1 also shows that the international distribution of disasters, combined with a risk aversion of only 4 implies an equity premium relative to risky government debt of 4.4%, close to the target of 4.8%. The equity premium that would be observed in samples without disasters is 4.7%. Most of the model's ability to explain the equity premium comes from the higher required return in the presence of disasters, as opposed to the bias in observations of returns over a sample without disasters. To summarize: disaster risk can explain the equity premium and riskfree rate puzzles, even if risk aversion is as low as 4, and even if default on government bills is taken into account.

3.6 Disaster probabilities and prices

We now turn to the question of how the disaster probability affects stock prices. Because this model is iid and the disaster probability is constant, we answer this question using comparative statics. The comparative statics results nonetheless lay the groundwork for the dynamic results to come.

The price-dividend ratio is given by

$$\begin{aligned} \frac{S_t}{D_t} &= E_t \int_t^\infty \frac{\pi_s D_s}{\pi_t D_t} ds & (18) \\ &= \left(\beta - \mu_D + \frac{1}{\psi} \mu - \frac{1}{2} \left(\gamma + \frac{\gamma}{\psi} - 2\phi\gamma \right) \sigma^2 \right. \\ &\quad \left. + \lambda E_\nu \left[\underbrace{\left(\left(1 - \frac{1}{\theta} \right) (e^{(1-\gamma)Z_t} - 1) - (e^{(\phi-\gamma)Z_t} - 1) \right)}_{\text{Effect of disaster probability on DP ratio}} \right] \right)^{-1}, & (19) \end{aligned}$$

where the last line is shown in Section A of the online appendix. It is the term multiplying the disaster probability, labeled “effect of disaster probability on DP ratio” that interests us. As in Campbell and Shiller (1988), we can decompose this term as follows:

$$\begin{aligned} \text{Effect of disaster probability on DP ratio} = & \\ & \underbrace{E_\nu \left[-(e^{-\gamma Z_t} - 1) + \left(1 - \frac{1}{\theta} \right) (e^{(1-\gamma)Z_t} - 1) \right]}_{\text{risk-free rate}} + \underbrace{E_\nu [(e^{-\gamma Z_t} - 1)(1 - e^{\phi Z_t})]}_{\text{equity premium}} \\ & - \underbrace{E_\nu [e^{\phi Z_t} - 1]}_{\text{expected dividend growth}}. \quad (20) \end{aligned}$$

Note that the terms labeled risk-free rate and equity premium are taken from (11) and (10) respectively. The last term is the direct effect of a disaster on the cash flows to equity.

Intuitively, an increase in the risk of a rare disaster should lower prices because it raises the equity premium and lowers expected cash flows. Equation 20 shows that there is an opposing effect coming from the riskfree rate. An increase in the risk of a disaster lowers

the riskfree rate, raising the price of any asset that is a store of value, including equities. Thus the net effect on the dividend-price ratio depends on which is greater: the riskfree rate effect or the (total) equity premium and cash flow effect. The answer is complicated in that it depends on γ , ψ and ϕ . We consider three special cases of interest to the literature. To fix ideas, assume that risk aversion γ is greater than one.

No leverage

This is the case we have been considering thus far in this chapter, and it has the appeal of parsimony, since we have not yet introduced a theory for why dividends would respond more than consumption in the event of a disaster. The term inside the expectation in (19) is positive if and only if $\theta < 0$, namely if and only if the EIS, ψ , is greater than 1.

Time-additive utility

In this case, $\theta = 1$. The term inside the expectation is positive if and only if $\phi > \gamma$, namely the responsiveness of dividends in the event of a disaster (or “leverage”) exceeds risk aversion.

EIS = 1

In this case, $1 - 1/\theta = 1$. The term inside the expectation is positive if and only if $\phi > 1$, namely if and only if dividends are more responsive to disasters than consumption

Note that, while we use the same parameter ϕ to determine the responsiveness of dividends to disasters and to normal shocks, it is in fact only the responsiveness of dividends to disasters that determines the direction of the effect. Note too that this effect is independent of the debate concerning government default. This changes the decomposition of the discount rate into a government bill rate (higher than the riskfree rate) and an equity premium relative to the government bill rate (lower than the equity premium relative to the riskfree rate). It does not change the total discount rate, which is what matters for the discussion above.

Generally, the higher is the EIS, the less the riskfree rate response to a change in the

probability, and the lower the precautionary motive. The greater the responsiveness of dividends, the greater the equity premium and cash flow effect combined. Section 3.7 and Section 4 introduce dynamic models which differ in their details. However, in each of these models, prices are subject to the same simple economic forces outlined here.

What happens empirically when the probability of a disaster increases? Using a political science database developed for the purpose of measuring the probability of political crises, Berkman, Jacobsen, and Lee (2011) show that an increase in the probability of a political crisis has a large negative effect on world stock returns, with the effect size increasing in the severity of the crisis. Thus the data clearly favor parameter values that imply a negative relation between prices and the risk of a disaster.

3.7 Multiperiod disasters

We now return to a question raised in the previous section. In the model of Section 3.1, disasters occur instantaneously. In the data, they unfold over several years.¹³ How does this affect the model's ability to explain the equity premium?

One simple way to model multiperiod disasters is to assume disasters are in expected rather than realized consumption.

$$\frac{dC_t}{C_t} = \mu_t dt + \sigma dB_t \quad (21)$$

with

$$d\mu_t = \kappa_\mu(\bar{\mu} - \mu_t) dt + Z_t dN_t. \quad (22)$$

In this model, consumption itself is continuous, and the normal distribution describes risks over infinitesimal time periods. However, over a time period of any finite length, consumption growth will exhibit fat tails. If the model in the previous section represents one extreme, this represents another. Most likely reality is somewhere between the two in that some part of

¹³The average duration of the disasters in Barro and Ursúa (2008), measured as the number of years from consumption peak to trough, is 4.3 years. The size and duration of the disaster has a negative correlation of -0.40.

the disaster is instantaneous, while another part is predictable.¹⁴ Economically, this model seems quite similar to the one in Section 3.1, and we will see that basic similarity does carry through to the conclusions, with some twists.

An important statistic in this model is the cumulative effect of a disaster on consumption, equal to Z_t/κ_μ (see Appendix C). Longer disasters can be captured by lowering κ_μ ; to match the consumption data one would then lower Z_t to keep Z_t/κ_μ the same. Asset pricing results remain largely unaffected by changes in the calibration that preserve the key quantity Z_t/κ_μ .¹⁵

The state-price density in this model has an approximate analytical solution that is exact when the EIS is equal to one and (in a trivial sense) when utility is time-additive:¹⁶

$$\frac{d\pi_t}{\pi_t} = \mu_\pi dt - \gamma\sigma dB_t + \left(e^{(i_1 + \kappa_\mu)^{-1}(\frac{1}{\psi} - \gamma)Z_t} - 1 \right) dN_t. \quad (23)$$

Here, $i_1 = e^{E[c-w]}$, where $c - w$ is the log consumption-wealth ratio in the economy. When the EIS is equal to 1, $i_1 = \beta$. Under time-additive utility, $\frac{1}{\psi} = \gamma$ and thus the shocks to the state-price density are the same as if there were no disasters. In a typical calibration, where disasters unfold over years rather than decades, κ_μ is two orders of magnitude greater than i_1 . Thus the term multiplying the Poisson shock is well-approximated by $e^{(\frac{1}{\psi} - \gamma)\frac{Z_t}{\kappa_\mu}} - 1$. This term, which determines the risk premium for bearing disaster risk, is approximately invariant to changes in the consumption process that leave Z_t/κ_μ unchanged.

¹⁴An alternative way to produce disaster clustering is to assume that the disaster probability follows a self-exciting process, in that an occurrence of a disaster raises the probability of future disasters. Such a model is solved in closed form by Nowotny (2011). It is also the case that, even if such a clustering is not a property of the physical process of disasters (as it appears to be), it occurs endogenously through learning, as shown by Gillman, Kejak, and Pakoš (2014). Multifrequency processes can produce dynamics that resemble disasters (Calvet and Fisher (2007)). One can also assume that, rather than a one-time event, a disaster represents a state in which there is some probability of entry and exit as in Nakamura, Steinsson, Barro, and Ursúa (2013). The conclusions we draw are robust to alternative specifications.

¹⁵This statement holds provided that κ_μ is “large” in a sense that will be made more precise as we go along. In effect what is required is that we look at disasters that last for several years rather than several decades.

¹⁶See Tsai and Wachter (2014a) for details.

To solve for the equity premium, we first find the process for stock prices:

$$\frac{S_t}{D_t} = E_t \int_t^\infty \frac{\pi_s}{\pi_t} \frac{D_s}{D_t} ds. \quad (24)$$

We can solve the expectation to find

$$\frac{S_t}{D_t} = G(\mu_t) = \int_0^\infty e^{a_\phi(\tau) + b_\phi(\tau)\mu_t} d\tau,$$

where $b_\phi(\tau)$ takes a particularly simple form:

$$b_\phi(\tau) = \frac{\phi - \frac{1}{\psi}}{\kappa_\mu} (1 - e^{-\kappa_\mu \tau}). \quad (25)$$

In fact, because κ_μ is on the order of 1, $b_\phi(\tau) \approx \frac{\phi - \frac{1}{\psi}}{\kappa_\mu}$, which does not depend on τ . Therefore, from Proposition 1 it follows that the equity premium can be approximated by

$$r^S - r \approx \gamma\phi\sigma^2 - \lambda E_\nu \left[(e^{(\frac{1}{\psi} - \gamma)\frac{Z_t}{\kappa_\mu}} - 1)(e^{(\phi - \frac{1}{\psi})\frac{Z_t}{\kappa_\mu}} - 1) \right] \quad (26)$$

See Appendix C for the full solution to the model. Note the similarity to the equity premium in the model we previously considered (keeping in mind that Z_t/κ_μ represents the total size of the disaster and is comparable to Z_t in that model). The difference is that in the disaster premium term, risk aversion is replaced by the difference between risk aversion and the inverse of the EIS $\gamma - \frac{1}{\psi}$, and leverage is replaced by the difference between leverage and the inverse of the EIS: $\phi - \frac{1}{\psi}$. The presence of the EIS in these terms reflects the response of the riskfree rate when disasters are not instantaneous. Upon the onset of a multiperiod disaster, the riskfree rate falls. This offsets the effect of the decline in expected cash flows.

Table 2 shows the values of the equity premium and for the observed excess returns, assuming risk aversion equal to 4, and three cases for the EIS: time-additive utility (corresponding to an EIS of 1/4 in this case), unit EIS, and EIS equal to 2. We show results for the consumption claim ($\phi = 1$) and for a levered claim ($\phi = 3$). We simplify our calibration

by considering $\kappa_\mu = 1$, which approximately matches the duration of disasters in the data. Thus Z_t remains the cumulative effect of disasters in this model, just as in Section 3.1. Time-additive utility implies a population equity premium that is the same as the CCAPM; disasters do not contribute anything in this model. The observed average excess return in a sample without disasters will actually be negative. That is because equity prices rise on the onset of a disaster in the time-additive utility model (with leverage below risk aversion). Investors factor this price increase into their equity premium, and thus when it does not occur, the observed return is in fact lower. In contrast, for EIS equal to 1, the premium on the consumption claim is the same as in the CCAPM in samples with and without disasters. For an unlevered claim with EIS equal to 1, there is no price change when a disaster occurs, and no risk premium (other than for normal Brownian risk). For a reasonable value of leverage ($\phi = 3$), the model implies an equity premium of 5.15% in samples without disasters when $\psi = 1$. When $\psi = 2$, this equity premium is 8.05%. In the data, the equity premium is 7.69%. To conclude, the model can still explain the equity premium puzzle, even if disasters are spread over multiple periods.

4 Time-varying Risk Premia

The previous sections focused on the question of whether agents' beliefs about rare disasters can explain the equity premium puzzle. Another important puzzle in asset pricing is the high level of stock market volatility. The basic endowment consumption CAPM fails to explain this volatility: assuming (for the moment) that stock returns are not levered, consumption volatility and return volatility should be equal. However, in postwar data consumption growth volatility is 1.3%, while stock return volatility is 18%. Adding leverage helps a little bit, but not very much. Dividends are more volatile than consumption, though not be nearly enough to account for the difference between volatility of consumption and volatility of returns. This striking difference between the volatility of cash flows and the volatility of returns is known as the *equity volatility puzzle* (Shiller (1981), LeRoy and Porter (1981),

Campbell and Shiller (1988)).¹⁷ Closely connected with the volatility puzzle is the fact that excess stock returns are predictable by the price-dividend ratio (Cochrane (1992), Fama and French (1989), Keim and Stambaugh (1986) and many subsequent studies). The reason is that, in a rational, stationary, model, returns are driven by realized cash flows, expected future cash flows, and discount rates. Empirical evidence point strongly to discount rates, and in particular, equity premia, as being the source of this variation; and if equity premia vary, this should be apparent in a predictive relation between the price-dividend ratio and returns (Campbell (2003)).

The models shown so far do not help explain the equity volatility puzzle. In the iid model of Section 3.1, disaster volatility and consumption volatility (or dividend volatility in the levered model) are the same. Depending on the calibration, this model might have a much greater volatility of stock returns than the consumption CAPM. However, this volatility arises entirely from the behavior of stock returns and consumption during disasters. It says nothing about why stock return volatility would be high in periods where no consumption disaster has occurred. In the multiperiod disaster model of Section 3.7, stock prices fall by more than consumption upon the onset of a disaster. In this sense, the model does have some excess volatility. However, like the iid model, it cannot explain the volatility of stock returns during normal times.¹⁸

To fully account for the evidence discussed above, the source of variation in returns should be through risk premia. The disaster risk framework offers a natural mechanism

¹⁷Some have pointed out that total payouts to stockholders are themselves very volatile, and that this may be a better measure of cash flows than dividends (Boudoukh, Michaely, Richardson, and Roberts (2007), Larrain and Yogo (2008)). While interesting, this fact does not lead to a solution to the volatility puzzle. A model should be able to explain both the behavior of the “dividend claim” and the behavior of the “cash flow claim”, as both represent cash flows from a replicable portfolio strategy. Moreover, as explained further above, excess returns are predictable by price-dividend ratios, implying that part of return volatility comes from time-varying equity premia. Price-to-cash flow measures imply even greater evidence of return predictability than the traditional price-dividend ratio.

¹⁸Time-variation in consumption growth can be used to explain the equity premium and normal-times equity volatility (Bansal and Yaron (2004)). However, this mechanism, by itself, does not explain return predictability, and it implies that consumption growth is predictable by the price-dividend ratio, which it does not appear to be (Beeler and Campbell (2012)). The multiperiod disaster model in Section 3.7 also implies that consumption growth is predictable by the price-dividend ratio, but only in samples in which a disaster takes place.

by which this can occur. Rather than being constant, the probability of a disaster might vary over time. Consider (10): the equity premium depends directly on the probability of a disaster. Consider also the discussion in Section 3.6. Under reasonable parameter specifications, an increase in the disaster probability leads to a decline in the price-dividend ratio. This reasoning suggests that such a model could account for stock market volatility through variation in the disaster probability, as well as for the predictability in stock returns.

4.1 Time-varying probability of a disaster

A dynamic model that captures this intuition is described in Wachter (2013). Consumption and dividends are as in Section 3.1, except that here the probability of a disaster is time-varying:

$$d\lambda_t = \kappa_\lambda(\bar{\lambda} - \lambda_t) dt + \sigma_\lambda \sqrt{\lambda_t} dB_{\lambda,t}, \quad (27)$$

where $B_{\lambda,t}$ is also a standard Brownian motion, assumed (for analytical convenience) to be independent of B_t . The distribution of Z_t is also assumed to not depend on λ_t . This process has the desirable property that λ never falls below zero; the fact that it is technically an intensity rather than a probability implies that the probability of a disaster can never exceed one.¹⁹

Wachter (2013) assumes a unit EIS. Here, we generalize to any positive EIS; as in Section 3.7 our solutions are exact in the time-additive and unit EIS case and approximate otherwise. As we show in Tsai and Wachter (2014a), the state-price density for this model is

$$\frac{d\pi_t}{\pi_{t-}} = \mu_\pi dt - \gamma\sigma dB_t + \left(1 - \frac{1}{\theta}\right) b\sigma_\lambda \sqrt{\lambda_t} dB_{\lambda,t} + (e^{-\gamma Z_t} - 1) dN_t, \quad (28)$$

where b is an endogenous preference-related parameter given by

$$b = \frac{\kappa_\lambda + i_1}{\sigma_\lambda^2} - \sqrt{\left(\frac{\kappa_\lambda + i_1}{\sigma_\lambda^2}\right)^2 - 2\frac{E_\nu[e^{(1-\gamma)Z_t} - 1]}{\sigma_\lambda^2}}, \quad (29)$$

¹⁹As in Wachter (2013), we abstract from the multiperiod nature of disasters described in Section 3.7. In other work (Tsai and Wachter (2014b)) we consider multiperiod disasters that occur with time-varying probability.

where $i_1 = e^{E[c-w]}$, equal to β in the case of unit EIS. Note that when $\gamma > 1$, $b > 0$. Note that if utility is time-additive ($\theta = 1$), the pricing kernel does not depend on b . Otherwise, shocks to λ_t are priced, even though they are uncorrelated with shocks to consumption. The agent's preference for the early resolution of uncertainty determines how shocks to the disaster probability are priced, just as it determines pricing of the multiperiod disaster shocks in Section 3.7. For example, consider $\gamma > 1$. Then a preference for early resolution of uncertainty ($\gamma > 1/\psi$) implies $1 - \frac{1}{\theta} > 0$. Assets that go down in price when the disaster probability rises require an additional risk premium to compensate for disaster probability risk. The case of $\gamma < 1$ works in a similar fashion.

The riskfree rate in this economy is equal to

$$r_t = \underbrace{\beta + \frac{1}{\psi}\mu - \frac{1}{2}\gamma\left(1 + \frac{1}{\psi}\right)\sigma^2}_{\text{CCAPM}} - \underbrace{\frac{1}{2}\frac{1}{\theta}\left(\frac{1}{\theta} - 1\right)b^2\sigma_\lambda^2\lambda_t + \lambda_t E_\nu \left[\underbrace{\left(1 - \frac{1}{\theta}\right) \left(e^{(1-\gamma)Z_t} - 1\right) - \left(e^{-\gamma Z_t} - 1\right)}_{\text{constant disaster risk}} \right]}_{\text{time-varying disaster risk}}. \quad (30)$$

Relative to the iid model (11) in Section 3.4, there is an additional term that reflects precautionary savings due to time-variation in λ_t itself. This term is zero for both time-additive utility and unit EIS (notice it is multiplied by both $\frac{1}{\theta}$ and $\frac{1}{\theta} - 1$). If there is a preference for an early resolution of uncertainty, and if $\gamma > 1$ and $\psi > 1$, then $\theta < 0$ and this term lowers the riskfree rate relative to what it would be in a model with constant disaster risk.

We now turn to equity pricing. The price-dividend ratio satisfies the general pricing equation (24). Solving the expectation (see Tsai and Wachter (2014a) for details results in a similar form to that in (3.7)):

$$\frac{S_t}{D_t} = G(\lambda_t) = \int_0^\infty e^{a_\phi(\tau) + b_\phi(\tau)\lambda_t} d\tau \quad (31)$$

where $a_\phi(\tau)$ and $b_\phi(\tau)$ are functions available in closed form (see Appendix D for more

details). As in Section 3.7, the value of the aggregate market is the integral of the values of “zero-coupon equity” claims, or equity “strips”, namely assets that pay the aggregate dividend at a specific point in time and at no other time. For a zero-coupon claim of maturity τ , $b_\phi(\tau)$ determines how that claim responds to a change in the disaster probability. The same tradeoffs that govern the comparative statics results in Section 3.6 determine how prices respond to changes in disaster probabilities, as explained in Appendix D. We will focus attention on the cases that lead prices to decline in the disaster probability, namely such that the risk premium and expected cash flow effects outweigh the riskfree rate effect.

Allowing the probability of the disaster to vary has implications both for the average (long-run) equity premium, and for the variability of the equity premium. Putting (31) together with (28) and Proposition 1 implies that the equity premium is given by

$$r_t^e - r_t = \underbrace{\phi\gamma\sigma^2}_{\text{CCAPM}} - \underbrace{\lambda_t \frac{G'(\lambda_t)}{G(\lambda_t)} \left(1 - \frac{1}{\theta}\right) b\sigma_\lambda^2}_{\text{time-varying disaster risk}} + \underbrace{\lambda_t E_\nu [(e^{-\gamma Z_t} - 1)(1 - e^{\phi Z_t})]}_{\text{constant disaster risk}}. \quad (32)$$

The expression for the equity premium has two components from the iid model (10), plus a term that reflects the compensation for disaster probability risk. $(1 - \frac{1}{\theta})b$ is positive in the case where there is a preference for early resolution of uncertainty. Because we are considering preference parameters such that the aggregate market goes down in price when the disaster probability rises, then the equity premium will be higher than in the iid model of Section 3.1.

To quantitatively assess the model, we simulate 600,000 years of data to obtain population moments, and also 100,000 samples of 65 years to represent the postwar data. We separately report statistics for those samples that do not contain disasters as these are the relevant comparisons. Table 3 shows the expected bond return, bond return volatility, the equity premium, stock return volatility, and the Sharpe ratio on the market for six calibrations: We consider three values of the EIS ($\psi = 1/3, 1$ and 2) and for each of these cases assuming both constant and time-varying disaster probability. We set risk aversion $\gamma = 3$; thus $\psi = 1/3$

refers to time-additive utility. Parameter values are reported in the caption of Table 3.²⁰

When $\psi = 1$ and 2, and when disaster risk is time-varying, the model is capable of matching the high equity premium, high equity volatility and a low and smooth Treasury Bill rate in the data.²¹ To understand where these results come from, we first discuss the constant disaster risk case and the time-additive utility case for comparison.

The constant disaster risk case is the model discussed in Section 3.1, though the parameter values are slightly different (specifically, risk aversion and the discount rate are lower, and there is leverage). We see that this model is capable of generating a low Treasury Bill rate and a reasonable equity premium. However, the model cannot generate reasonable stock return volatility because the volatility of stock returns is the same as the volatility of dividends.

Now consider the time-varying disaster risk case under time-additive utility. Even though disaster risk varies over time, equity returns are no more volatile in this case than in the constant disaster risk case, and in fact the equity premium is slightly lower. Equity returns are not more volatile because, under this calibration (leverage equals risk aversion), the riskfree rate, risk premium, and cash flow effects exactly cancel. Even if risk aversion were

²⁰Details of the calibration are as follows: normal-times consumption parameters μ_c and σ_c are set to match the postwar data. Leverage ϕ is set to 3, a standard value in the literature (Bansal and Yaron (2004)). Rather than assuming $D_t = C_t^\phi$, we follow Bansal and Yaron in assuming the more realistic specification

$$\frac{dD_t}{D_t} = \mu_D dt + \phi\sigma dB_t + \sigma_i dB_{it}$$

where B_{it} is a Brownian motion independent of B_t . We chose $\mu_D = 0.04$ to match the average price-dividend ratio in the data (note that share repurchases suggest expected dividend growth in the data may not accurately measure investors long-term expectations of dividend growth). We chose “idiosyncratic” dividend volatility $\sigma_i = 0.05$ so that measured dividend volatility in no-disaster samples in the model equals dividend volatility in postwar U.S. data. We assume that there is a 40% probability of default on government bills in the case of a disaster. The time-varying disaster risk parameters κ_λ and σ_λ are of course not directly observable. The existence of the likelihood function imposes tight constraints on these parameters so in practice there is a single free parameter that must do its best to match both the volatility of stock returns, the persistence of the price-dividend ratio, and the volatility of Treasury Bill returns. See Section B of the online Appendix for more details.

²¹One concern about high values of the EIS is that generate a counterfactual relation between the riskfree rate and consumption growth in Hall (1988) regressions (see the discussion in Beeler and Campbell (2012)). These regressions use two stage least squares to identify the EIS, and thus require some predictability in consumption growth. In the model above, there is no predictability in consumption growth and thus running regressions in samples without disasters will replicate the data finding of a small and insignificant EIS. Of course, this should not be taken as license to make the EIS arbitrarily high, since such a model might deliver counterfactual predictions in a realistic extension with a small amount of consumption growth predictability.

to slightly exceed leverage, equity returns would still not be sufficiently volatile, and if risk aversion were lower than leverage, the model would generate return predictability of the incorrect sign. At first glance it is surprising that the measured equity premium is slightly lower in no-disaster simulations in the time-varying case. After all, in the constant case, in the no-disaster simulations, the measured equity premium is lower because disasters have not in fact been realized. In the time-varying case, this result is offset by a second sample selection problem. Samples without disasters tend to have a lower probability of disaster and hence, in this model, a lower equity premium.

Finally consider the time-varying case when the EIS is greater than or equal to one. In these cases, equity volatility is greater in the time-varying disaster case because the risk premium (and cash flow) effect exceeds the riskfree rate effect. Indeed, the Treasury Bill return is less volatile in these cases than in the time-additive utility case. Meanwhile, the equity premium is higher because of the extra compensation required for assets that fall in times when the disaster probability is high.

Given that the mechanism is through a time-varying risk premium, the model also explains the ability of price-dividend ratios to predict future excess stock returns (Wachter (2013)). Moreover, consumption growth will not be predictable, as in the data. One limitation of this model is that it does not predict as much volatility of the price-dividend ratio as in the data. The model can come close to matching this volatility in population, and in the full set of simulations, but the volatility is lower (0.21 vs 0.43 in the data) in the simulations without rare disasters. This in part arises from the particular functional form imposed above, which imposes tight restrictions on the volatility of the disaster process. Richer models for the disaster probability process can overcome these problems (Seo and Wachter (2015)).

The key claim of this model is that asset prices during normal times are driven by the time-varying tail behavior. Investor expectations about this tail behavior are of course hard to capture, but recent work has made progress in this regard. Bollerslev and Todorov (2011) use option prices and high-frequency data to determine investor beliefs about tail risk, and premia investors assign to tail events. Their model-free measures find that the risk premia

that can be attributed to tail events is a large fraction of the total equity premium, and that this premium is strongly varying over time. Kelly and Jiang (2014) assume that tails of individual asset returns follow a power law, with parameters determined by overall market tail behavior. This structure allows them to estimate the probability of an overall market disaster from individual stock returns. They show that their disaster measure captures risk premia, in that it has strong predictive power for excess stock returns. They also show that tail risk determines the cross-section in that stocks that covary more with the their time-varying tail risk measure have high expected returns. Chabi-Yo, Ruenzi, and Weigert (2014) use high-frequency data to measure disaster risk and come to similar conclusions, as do Gao and Song (2013) and Gao, Song, and Yang (2014) who use data on a wide cross-section of options to construct a disaster-risk measure. Disaster risk can also be measured by news articles using text-based analysis (Manela and Moreira (2013)) and using data on credit-default swaps (Seo (2014)). These empirical findings offer direct confirmation of the mechanism in the model above. A related finding is that investors' lifetime experience of disasters affects their allocation of stocks years later (Malmendier and Nagel (2011)). Though technically outside of this representative-agent framework, this finding is very much in the spirit of the model in that it shows the strong hold that rare events can have on investor attitudes.²²

4.2 Time-varying Resilience

An alternative mechanism for generating time-varying risk premia is to allow the sensitivity of cash flows to a disaster to vary. This mechanism is introduced by Gabaix (2012), who makes use of linearity-generating processes, defined in Gabaix (2008), to obtain simple and elegant solutions for prices.

Following Gabaix (2012), consider the discrete-time equivalent of (5), except with the

²²One study that takes heterogeneity in beliefs into account is Chen, Joslin, and Tran (2012). Challenges in modeling disagreements about rare events include counterparty risk: namely, can optimistic agents commit to insuring pessimistic ones in the event of a disaster, and the well-known problem in heterogeneous-agent models of non-stationary equilibria (optimistic agents, over time, take over the economy).

Gaussian shock set to zero:

$$\log \frac{C_{t+1}}{C_t} = \mu + \begin{cases} 0 & \text{if there is no disaster at } t+1 \\ Z_{t+1} & \text{if there is a disaster at } t+1 \end{cases} \quad (33)$$

Assume the following process for dividends:

$$\log \frac{D_{t+1}}{D_t} = \mu_D + \log(1 + \epsilon_{t+1}^D) + \begin{cases} 0 & \text{if there is no disaster at } t+1 \\ Z_{d,t+1} & \text{if there is a disaster at } t+1 \end{cases} \quad (34)$$

where $Z_{d,t+1}$ is the response of dividends to the disaster and $\epsilon_{t+1}^D > -1$ has mean zero and is independent of the disaster event.²³ Assume that the pricing kernel is given by:

$$\log \frac{\pi_{t+1}}{\pi_t} = -\beta - \gamma \log \left(\frac{C_{t+1}}{C_t} \right)$$

which is equivalent to assuming time-additive utility. The distribution for $Z_{d,t+1}$ and the disaster intensity are jointly specified using a quantity H_t , called resilience:

$$H_t = \lambda_t E_\nu \left[e^{-\gamma Z_{t+1} + Z_{d,t+1}} - 1 \right] \quad (35)$$

where λ_t is the disaster probability and the expectation is taken over the disaster distributions. The process for H_t is given by

$$H_t = H_* + \hat{H}_t, \quad (36)$$

where

$$\hat{H}_{t+1} = \frac{1 + H_*}{1 + \hat{H}_t} e^{-\phi_H} \hat{H}_t + \epsilon_{t+1}^H, \quad (37)$$

where ϵ_{t+1}^H is mean zero and uncorrelated with ϵ_{t+1}^D and the disaster event. These assumptions

²³Gabaix (2012) allows for the possibility that dividends are completely wiped out in a disaster; for notational convenience we can allow for this possibility by formally setting $e^{Z_d} = 0$.

imply that the price-dividend ratio on the market is linear in the state variable H_t :

$$\frac{S_t}{D_t} = \frac{1}{1 - e^{-\delta}} \left(1 + \frac{e^{-\delta - h_*} \hat{H}_t}{1 - e^{-\delta - \phi_H}} \right), \quad (38)$$

where $h_* = \log(1 + H_*)$ and $\delta = \beta - \mu_D - h_*$. Assumptions (34–37) are designed to ensure this linearity. Comparing this model with the previous one, one sees an example of a general principle, which is that there is a trade-off in the complexity of assumptions versus the complexity of the conclusions. The model in Section 4.1 has relatively simple assumptions on economic fundamentals, and results in prices that are an integral of exponential functions. This model has more complex processes for fundamentals but implies that prices are linear functions.²⁴

Like the model in Section 4.1, this model also implies that risk premia vary over time, and thus that excess stock returns are predictable. The mechanism, however, is different. Gabaix (2012) focuses on the case with constant probability in (35), so the main calibration of his model is equivalent to one with time-varying exposure to a disaster. Low market values correspond to periods of high exposure to a constant level of disaster risk. Because disaster exposure is high, risk premia are high. Thus excess stock returns are predictable as in the data. The reason for this focus, perhaps, is the assumption of time-additive utility. Under this assumption, variation in the probability of a disaster will help little in matching the data, and may hurt, because the riskfree rate effect will be large in comparison to the risk premium effect and cash flow effect.²⁵

²⁴Which one is preferable as a modeling strategy will likely depend on the application. One issue is the utility function: while Gabaix (2012) shows that linearity-generating processes can be generalized to non-time-additive utility functions, this comes at the cost of greater complexity. A second question is whether evaluating the model involves simulation and what quantities are necessary to calculate; it is more convenient to calculate prices as linear functions, but more convenient to simulate from the model in Section 4.1. Stability and the absence of arbitrage require conditions on the shocks ϵ_{t+1}^H and ϵ_{t+1}^D . Provided that these conditions are met, the distribution of prices will not be affected by the distributional assumptions on these shocks. However, any quantity that is not a price (such as a return) will be affected. Thus to simulate quantities other than prices from this model one must take a stand on the distribution for these shocks; simple distributional assumptions will generally not be sufficient because of the required conditions.

²⁵This can be seen directly from (35); an increase in the probability of a disaster will increase resilience if $E_\nu [e^{-\gamma Z_{t+1} + Z_{d,t+1}} - 1] > 0$, namely if risk aversion is high relative to leverage. This increase will *raise* prices, at the same time that it raises risk premia.

5 Alternative asset classes

The previous sections focused the implications of disaster risk on stock returns and short-term bonds, while briefly touching on the literature linking disaster risk to the cross-section. Here, we discuss models that account for puzzles in other asset classes.

5.1 Fixed income

Like excess stock returns, excess returns on long-term bonds over short-term bonds are also predictable (Campbell and Shiller (1991), Stambaugh (1988), Cochrane and Piazzesi (2005)). Moreover, excess bond returns are positive on average, leading to an upward-sloping yield curve. Gabaix (2012) and Tsai (2014) write down models explaining these results. Inflation, and thus low realized returns on bonds occur frequently in disasters. Thus, by the logic of Proposition 1, nominal bonds will carry a disaster premium because of their inflation exposure. If this premium varies over time because of time-varying inflation exposure (as Gabaix assumes) or because the probability of an inflation disaster varies (as Tsai assumes), then bond returns will be predictable. Moreover, as Tsai shows, bond and stock risk premia need not move together, because fears of an inflation disaster need not coincide with a disaster affecting stock returns. This lack of co-movement poses a puzzle for preference-based theories of time-varying risk premia (see discussions in Duffee (2013) and Lettau and Wachter (2011)).

Rare disasters can also help account for the puzzling size of the default spread, as compared to actual defaults, as shown in models of capital structure choice in the presence of rare disasters (Bhamra and Strebulaev (2011), Gourio (2013)).

5.2 Options

By making use of the tight connection between the Merton (1976) model and the iid disaster model in Section 3.1, Backus, Chernov, and Martin (2011) derive an elegant method for

backing out the distribution of consumption disasters implied by the prices of equity index options. Their method makes use of the fact that, in an iid model, stock returns and consumption returns move one-to-one with each other. Moreover, because of their non-linear payoff structure, options provide a means of determining the risk neutral distribution of stock returns. Finally, by assuming a representative agent with constant relative risk aversion and choosing risk aversion to match the equity premium, Backus et al. can create a mapping from the risk-neutral distribution to the physical distribution.

The distribution of consumption growth implied by options, according to this method, looks very different from the distribution assumed in Section 3.1. Specifically, Backus, Chernov, and Martin (2011) show that options imply “disasters” in consumption that are much smaller, and also much more likely. A normal distribution for consumption growth is still clearly rejected, and yet the distribution looks nothing like the distribution of macroeconomic disasters. Backus et al. note that their calibration can be seen as an alternative disaster calibration to that of Section 3.1: namely an alternative calibration of the same model that can also explain the equity premium, but, unlike the model in Section 3.1, can also explain option returns (for, just as the implied consumption distribution looks very different than the international macro distribution of disasters, the option prices implied by the model of Section 3.1 look very different from options in the data). One difficulty with this interpretation of their findings is that their consumption distribution is inconsistent with postwar consumption data. For the same reason consumption growth cannot be normally distributed with a volatility of 3.5%, given the 65 years of postwar data, it is also not possible for consumption to be subject to frequent, small “disasters”; such disasters would have been observed and are inconsistent with the measured volatility of consumption growth. One is thus forced to conclude that their interpretation of options data is a serious challenge to the model in Section 3.1.

However, these results do not necessarily pose a challenge to the models in Section 4.1.²⁶

²⁶Du (2011) also shows that an iid disaster risk model has difficulty accounting for options data, but a model with both disaster risk and external habit formation utility as in Campbell and Cochrane (1999) can account for these data.

Seo and Wachter (2015) solve for option prices in the model in Section 4.1 and in a multi-factor extension. They show that, contrary to the model in Section 3.1, these models can reconcile international macro data, and hence the disaster-risk explanation of the equity premium, with option prices. This is surprising because allowing disaster risk to vary changes conditional, but not unconditional moments of the consumption distribution. Previous work has shown that the implied volatility smile is mainly driven by unconditional moments. Why then does allowing disaster risk to vary make a difference?

The reason is that the model in Section 4.1 can explain the volatility of stock returns outside of disasters while the model in Section 3.1 does not. Unlike reduced-form models in which the volatility of stock returns is an exogenous input, the disaster-risk models produce volatility endogenously. The level as well as the slope of the implied volatility curve reflects the fact that stock returns have high normal-times volatility, as well as a non-lognormal distribution. Given smooth consumption and dividend data, this is not consistent with an iid model.

Finally, a large and growing literature assesses the impact of disaster risk on the carry trade in foreign currency and in currency option returns. Due to space constraints we will not discuss this literature here, but rather refer the reader to Verdelhan (2015).

6 Production economies

The endowment economies above jointly specify agent preferences and endowments. It is more realistic to consider consumption as a choice variable, however. Thus an important question in considering representative agent models is whether the hypothesized consumption (and dividend) process is consistent with the deeper notion of general equilibrium defined by a model with production. Considering production economies also offers the promise of explaining correlations between the stock market and real variables such as investment and unemployment.

In this section of our survey we will follow the literature and assume a discrete-time

model. The discrete-time equivalent of the utility function (B.1,B.2) is

$$V_t = \left[(1 - e^{-\beta})C_t^{1-\frac{1}{\psi}} + e^{-\beta} (E_t[V_{t+1}^{1-\gamma}])^{\frac{1-\frac{1}{\psi}}{1-\gamma}} \right]^{\frac{1}{1-\frac{1}{\psi}}}, \quad (39)$$

(Epstein and Zin (1989), Weil (1990)). Here, it is more convenient to define the stochastic discount factor rather than the state-price density:

$$M_{t+1} = e^{-\beta} \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{1}{\psi}} \left(\frac{V_{t+1}}{E_t[V_{t+1}^{1-\gamma}]^{\frac{1}{1-\gamma}}} \right)^{\frac{1}{\psi}-\gamma}, \quad (40)$$

so that for an asset with return R_{t+1} from t to $t + 1$, the agent's first-order condition (or equivalently, no-arbitrage) implies

$$E_t[M_{t+1}R_{t+1}] = 1. \quad (41)$$

6.1 Constant disaster risk models

Two models with endogenous consumption choice are considered by Barro (2009). In the first model, output is a function of productivity A_t and labor L_t :

$$Y_t = A_t L_t^\alpha, \quad (42)$$

where α is between 0 and 1. In this model, assumptions on the endowment process are replaced by assumptions on productivity, so that A_t is given by:

$$\log \frac{A_{t+1}}{A_t} = \mu + \epsilon_{t+1} + \begin{cases} 0 & \text{if there is no disaster at } t + 1 \\ Z_{t+1} & \text{if there is a disaster at } t + 1, \end{cases} \quad (43)$$

where ϵ_{t+1} is an iid normal shock. We modify (39) by replacing C_t with period utility that allows for preferences over leisure:

$$U_t = C_t(1 - L_t)^\chi. \quad (44)$$

This modification is not without consequence. When consumption is replaced by (44), ψ can no longer be interpreted as the EIS. However, the EIS is still greater than one if and only if $\psi > 1$ (Gourio (2012)).

In this model with no investment, consumption C_t is equal to output Y_t . The wage is equal to the marginal product of labor,

$$w_t = \alpha Y_t / L_t = \alpha C_t / L_t, \quad (45)$$

Setting marginal benefit equal to marginal cost leads to the labor supply function²⁷

$$L_t = 1 - \chi C_t / w_t \quad (46)$$

Moreover, from (45),

$$\frac{C_t}{w_t} = \frac{1}{\alpha} L_t.$$

Substituting into (46) implies that, in equilibrium $L_t = \frac{\alpha}{\alpha + \chi}$. Surprisingly perhaps, hours worked are constant, despite the fact that productivity fluctuates. A disaster leads to a fall in productivity, which leads to lower consumption and lower wages. These two effects exactly offset so that hours worked remain the same. Equation (42) implies that consumption behaves exactly as in the model of Section 3.1. Dividends are given by

$$D_t = Y_t - w_t L_t = Y_t - \alpha C_t = (1 - \alpha) C_t,$$

²⁷Specifically, taking partial derivatives in (44) leads to

$$C_t \chi (1 - L_t)^{\chi-1} = w_t (1 - L_t)^\chi.$$

Rearranging yields (46).

where we have substituted in from the labor supply function. Thus dividends are proportional to consumption, and also fall by the same percentage in a disaster. Furthermore, because L_t is a constant, the stochastic discount factor is still given by (40).

It follows from the Euler equation then that the equity premium in this model is exactly the same as the equity premium on the consumption claim in the model of Section 3.1. Moreover, because the consumption process and stochastic discount factor are the same, the riskfree rate will be the same as well. Thus the model's ability to explain the equity premium is preserved under this model with endogenous consumption.

Barro (2009) presents a second model with endogenous consumption, where, rather than labor choice, there is capital and investment. Output is given by

$$Y_t = AK_t, \tag{47}$$

where K_t is the capital stock. Unlike the model above, productivity, A is constant. Consumption satisfies

$$C_t = Y_t - I_t = AK_t - I_t,$$

where I_t is investment, and capital evolves according to

$$K_{t+1} = K_t + I_t - \delta_{t+1}K_t,$$

where δ_{t+1} is stochastic depreciation. In this model, disasters are represented by destruction of capital, modeled as a shock to depreciation:

$$\delta_t = \delta + \begin{cases} 0 & \text{if there is no disaster at } t \\ 1 - e^{Z_t} & \text{if there is a disaster at } t. \end{cases} \tag{48}$$

The link between disasters and capital destruction seems most direct in the case of wars. However, one could also interpret capital more broadly, as including intangibles such as employee or customer value. Moreover, economic crises may lead to re-allocations that make

certain kinds of capital worthless (see the discussion in Gourio (2012)).

Because output can be either consumed or costlessly reinvested, a standard formula applies for the gross return to capital:

$$R_{t+1}^S = 1 + A - \delta_{t+1}. \quad (49)$$

In this model, the expected return on unlevered equity is the same as the expected return on capital, and is determined by (49). Consumption, and thus investment, is determined by the Euler equation (41). Contrast this with the endogenous labor choice model, in which consumption is immediately determined from output, given that labor supplied is a constant, and the expected return on equity is then determined from the Euler equation.

In this model, investment is determined by the disaster probability. Let ζ be the fraction of capital that is invested. Based on the resource constraint $C_t = (A - \zeta)K_t$ and the Euler equation, one can show that

$$\zeta = \psi \left(\frac{\frac{1}{\psi} - 1}{\gamma - 1} \right) \lambda E [e^{(1-\gamma)Z_t} - 1] + \text{Terms that do not depend on } \lambda$$

(see the Online Appendix). Recall that $Z_t < 0$, so the term inside the expectation is positive when $\gamma > 1$ and negative when $\gamma < 1$. On the one hand, a greater probability of disaster lowers wealth and leads agents to invest more (the income effect). However, there is now a greater chance that any investment will be lost in a disaster; that depresses investment (the substitution effect). When $\psi > 1$, it is the latter effect dominates, and a greater probability of disaster lowers investment (in a comparative static sense). When $\psi < 1$, the former effect dominates.

Note that the constant investment ratio implies that C_t suffers the same percent decline as capital in the event of a disaster. It follows from (48) that the first-order implications of this model for consumption and therefore the riskfree rate are the same as in the iid endowment economy model of Section 3.1. Furthermore, the first-order implications for the equity premium because of the equation for the capital return (49).

What happens when there is both capital and labor? Gabaix (2011) and Gourio (2012) consider the following model for output:

$$Y_t = K_t^\alpha (A_t L_t)^{1-\alpha}. \quad (50)$$

If a disaster affects capital K_t and productivity A_t by the same amount, then the risk premium on the consumption claim is the same as that in the endowment economy with constant disaster risk. Consumption and the riskfree rate are also the same (also, hours worked do not change in a disaster). During normal times, macro quantities follow the same dynamics as in a model without disasters, though the level of investment and consumption relative to capital will be different.

These results demonstrate that the insights of the iid model in Section 3.1 are not dependent on the assumption of the endowment economy. That is, the basic mechanism linking disasters to the equity premium survives a model of endogenous consumption; in this sense this explanation is more robust than competing explanations of the equity premium (Kaltenbrunner and Lochstoer (2010), Lettau and Uhlig (2000)).

6.2 Time-varying disaster risk models

As Section 4 shows, many interesting questions in asset pricing lie outside the reach of iid models. Can the insights of dynamic endowment economies be extended to models with production? The literature on dynamic production economies without disasters (e.g. Jermann (1998), Boldrin, Christiano, and Fisher (2001), Kaltenbrunner and Lochstoer (2010)) illustrates the many challenges that this research agenda faces. For example, mechanisms that lead to procyclical investment lead to countercyclical dividends, and thus imply a very low equity premium. Indeed, one of the reasons that constant disaster risk models with production in Section 6.1 are successful in explaining the equity premium is that they do not rely on dynamics in productivity or preferences to generate an equity premium. Yet extending the insights of the dynamic endowment economies to production economies is a

crucial step in building our understanding of links between finance and the macroeconomy, and in obtaining a deeper understanding of financial markets themselves.

Gourio (2012) presents a production-based model with adjustment costs in which the probability of a rare disaster varies over time, as in Wachter (2013). In Gourio's model, output is still determined by (50), with disasters still having the same proportional affect on capital and productivity. Gourio also assumes an EIS of 2 and a risk aversion of 3.8. Because the EIS is greater than one, the substitution effect will dominate the income effect, and, because of the preference for early resolution of uncertainty, assets exposed to the risk of disaster will carry an additional risk premium (Section 4.1).

The main implication of Gourio's (2012) model is that investment falls when the probability of a rare disaster rises. Gourio has two mechanisms that combine to achieve this result. First, investment declines as a fraction of output as described in the comparative statics analysis of Section 6.1. Second, adjustment costs create a wedge between the value of a unit of investment within the firm and outside the firm. This wedge is smaller when the disaster probability is high, so there is less incentive to invest in this case. Thus Gourio's model can explain several important puzzles in macroeconomics that are related to investment. That is, the model explains why investment and output are much more volatile than consumption, and why these variables are imperfectly correlated. Furthermore, equity prices decline when the probability of a disaster increases, as in the model of Section 4.1. Thus the model can explain empirical correlations between investment and stock prices.

One real-world aspect of recessions that is missing from the Gourio (2012) model is unemployment. The model does produce variation in hours worked. However, agents do not decrease hours when the probability of a disaster rises because they cannot find jobs. Rather it is because they increase consumption (because they decrease investment), and consumption and leisure are complements. While agents are not happy when the disaster probability rises (it is a high marginal utility state), they are working their optimal number of hours.²⁸

²⁸The model implication that consumption rises when the disaster probability falls also seems like it might be counterfactual. However, this is not obvious as the probability of a disaster is not observable to the

A second difficulty with the Gourio (2012) model lies in its implications for stock market volatility. When the disaster probability rises in this model, investment and profits fall, and because investment falls by more than profits, dividends. This increase in dividends offsets the decrease in prices due to time-variation in the disaster probability. Thus the volatility on the unlevered claim is very low – lower than the volatility of output itself (Gourio (2009)). For this reason, Gourio (2012) assumes a high degree of leverage: firms issue debt in proportion to capital stock and there is default in the event of disaster with a low recovery rate. The levered claim in the model is realistically volatile. However, the leverage required to achieve this result appears to be significantly higher than that found in the data (50% versus 30%). Moreover, it appears that the unlevered market also appears to have a significant amount of volatility in the data (Larrain and Yogo (2008)), casting doubt on the theory that stock market volatility is entirely driven by leverage.

Petrosky-Nadeau, Zhang, and Kuehn (2013) also present a production model with a time-varying probability of disasters. They assume a production function with labor only, as in the first model in Section 6.1 but with $\alpha = 1$. Unlike in that frictionless model, hiring is subject to search frictions as in Diamond (1982) and Mortensen and Pissarides (1994). Their model has the appealing feature that consumption disasters arise even though shocks to productivity are normally distributed. Thus their model provides a mechanism through which ordinary shocks can endogenously lead to disasters. Their model also explains the volatility of unemployment and job vacancies. Moreover, the model implies that labor market tightness is driven by the probability of a disaster, and thus should predict stock returns. This is born out in the data.

Three assumptions lie behind the ability of the Petrosky-Nadeau, Zhang, and Kuehn (2013) model to produce endogenous disasters. First, wages are high and rigid, so a decline in output leads to a greater-than-proportional decline in dividends (in contrast to the investment-based model of Gourio (2012)). Petrosky-Nadeau et al. further assume a high

econometrician, and Gaussian fluctuations in productivity generate realistic consumption and investment correlations. Moreover, while consumption rises on impact, it then falls back to its steady state. Gourio (2012) describes some modifications to the model that would reduce this positive effect.

separation rate, and a fixed component in the marginal cost of vacancies. In effect, it costs more to post a vacancy when labor conditions are slack and thus when output is low. As a result, firms do not hire, and output falls still more. Thus a sequence of small shocks to output can lead, through this mechanism, to a large drop in output which resembles a disaster.

While it is advantageous to have a model that explains disasters, the Petrosky-Nadeau, Zhang, and Kuehn (2013) model's reliance on these three assumptions is a limitation. In particular, the last two assumptions are non-standard. The assumption that posting a vacancy is more expensive in bad times than in good times does not appear to have empirical support, and is difficult to interpret economically. The assumption of a high separation rate is contradicted by the data (in the data, the separation rate is 3% rather than 5%). A second limitation of this model, like that of Gourio (2012), lies in the implications for stock market volatility. The volatility of stock returns in the model is 11%, compared to an unlevered volatility of 13%. However, this is in simulated data that includes disasters. As shown in Table 3, stock return volatility tends to be higher in samples with disasters than in samples without. This, along with the endogenous consumption distribution in their sample, raises the question of whether their model can account for the combination of low volatility of consumption and output and high volatility of stock returns that characterizes the postwar period.

The models discussed in this section incorporate a time-varying probability of disaster risk to help resolve deep puzzles in macroeconomics, such as the level of normal-times volatility in investment and in unemployment. While the level of stock market volatility remains a puzzle, the rate of progress shown by these papers suggests that changing expectations of disasters have an important role in our understanding of the stock market, the macroeconomy and the links between the two.

7 Conclusion

We argue in this review that disaster risk offers a coherent and parsimonious framework for understanding asset pricing puzzles. These include the equity premium, riskfree rate and volatility puzzles, as well as puzzles in fixed income and option markets. We also show that the disaster risk framework survives a transfer from endowment models to production models, which endogenize the consumption and dividend processes that are taken as given in the endowment framework.

One question that frequently arises in the setting of rare events is the appropriateness of the assumption of a rational expectations equilibrium. A rational expectations equilibrium assumes that agents know the probabilities with which all events can occur and can also forecast prices in all states of the world. This has always been somewhat of a heroic assumption in economics, and never more so in the case of rare events. Some authors relaxed this assumption (for example, Weitzman (2007), Collin-Dufresne, Johannes, and Lochstoer (2013)); a theme that emerges from this literature is that the uncertainty and learning strengthen the results that we report.

Disaster risk explanations involve a certain amount of technical modeling, but they are intuitive. For example, the equity premium puzzle comes down to a question about why investors do not desire to hold more equities; according to the disaster risk explanation it is because they are concerned about equity performance in what might be a second Great Depression. Another question is why time-varying risk premia (and thus excess volatility) persists, when investors can simply hold more equities in response to greater risk premia. According to the disaster risk explanation, this is when investors most fear disasters. Thus disaster risk is part of a broader class of models (see Hansen (2007)) whose goal is endowing agents with something like the uncertainty faced by actual agents making real-world decisions. While not alone in this goal, the rare disaster literature accomplishes this in a simple and powerful way.

Appendix

A Proof of Proposition 1

Proof of Proposition 1 S is the price of the asset that pays a continuous dividend stream D_t , by no arbitrage:

$$S_t = E_t \left[\int_t^\infty \frac{\pi_s}{\pi_t} D_s ds \right]. \quad (\text{A.1})$$

Multiplying each side of (A.1) by π_t implies

$$\pi_t S_t = E_t \left[\int_t^\infty \pi_u D_u ds \right]. \quad (\text{A.2})$$

The same equation must hold at any time $s > t$:

$$\pi_s S_s = E_t \left[\int_s^\infty \pi_u D_u ds \right]. \quad (\text{A.3})$$

Combining (A.2) and (A.3) implies

$$\pi_t S_t = E_t \left[\pi_s S_s + \int_t^s \pi_u D_u ds \right]. \quad (\text{A.4})$$

Adding $\int_0^t \pi_u D_u du$ to both sides of (A.4) implies

$$\pi_t S_t + \int_0^t \pi_u D_u du = E_t \left[\pi_s S_s + \int_0^s \pi_u D_u ds \right]. \quad (\text{A.5})$$

Therefore $\pi_t S_t + \int_0^t \pi_u D_u du$ is a martingale.

$$\begin{aligned} \pi_t S_t + \int_0^t \pi_u D_u du &= \pi_0 S_0 + \int_0^t \pi_u S_u \left(\mu_{\pi,u} + \mu_{S,u} + \frac{D_u}{S_u} + \sigma_{\pi,u} \sigma_{S,u}^\top \right) \\ &\quad + \int_0^t \pi_u S_u (\sigma_{S,u} + \sigma_{\pi,u}) dB_u + \sum_{0 < u_i \leq t} \left(\pi_{u_i} S_{u_i} - \pi_{u_i^-} S_{u_i^-} \right), \quad (\text{A.6}) \end{aligned}$$

where $u_i = \inf\{u : N_u = i\}$. Adding and subtracting the jump compensation term from (A.6) yields:

$$\begin{aligned} \pi_t S_t + \int_0^t \pi_u D_u du &= \pi_0 S_0 + \int_0^t \pi_u S_u \left(\mu_{S,u} + \mu_{\pi,u} + \frac{D_u}{S_u} + \sigma_{\pi,u} \sigma_{S,u}^\top + \lambda_u^\top E_\nu [(e^{Z_\pi + Z_S} - 1)] \right) ds \\ &+ \int_0^t \pi_u S_u (\sigma_{S,u} + \sigma_{\pi,u}) dB_u \\ &+ \left(\sum_{0 < u_i \leq t} (\pi_{u_i} S_{u_i} - \pi_{u_i^-} S_{s_i^-}) - \int_0^t \pi_u S_u \lambda_u E_\nu [(e^{Z_\pi + Z_S} - 1)] du \right). \end{aligned} \quad (\text{A.7})$$

The second and the third terms on the right-hand side of (A.7) are martingales. Therefore, the first term in (A.7) must also be a martingale, and it follows that the integrand of this term must equal zero:

$$\mu_{\pi,t} + \mu_{S,t} + \frac{D_t}{S_t} + \sigma_{\pi,t} \sigma_{S,t}^\top + \lambda_t^\top E_\nu [e^{Z_\pi + Z_S} - 1] = 0, \quad (\text{A.8})$$

It also follows from no-arbitrage that $\mu_{\pi,t} = -r_t - \lambda_t E_\nu [e^{Z_\pi} - 1]$, substituting this into (A.8) and re-arranging implies:

$$\mu_{S,t} + \frac{D_t}{S_t} - r_t = -\sigma_{\pi,t} \sigma_{S,t}^\top - \lambda_t (E_\nu [(e^{Z_\pi + Z_S} - 1)] - E_\nu [e^{Z_\pi} - 1]).$$

Add both sides by $\lambda_t E_\nu [e^{Z_S} - 1]$, and use the definition of r^S , we get

$$\begin{aligned} r_t^S - r_t &= -\sigma_{\pi,t} \sigma_{S,t}^\top - \lambda_t (E_\nu [(e^{Z_\pi + Z_S} - 1)] - E_\nu [e^{Z_\pi} - 1] - E_\nu [e^{Z_S} - 1]) \\ &= -\sigma_{\pi,t} \sigma_{S,t}^\top - \lambda_t E_\nu [(e^{Z_\pi} - 1)(e^{Z_S} - 1)] \end{aligned} \quad (\text{A.9})$$

□

B Utility

We specify the preferences of the representative agent using a non-time-additive generalization of constant relative risk aversion. The form of the utility function that we use is from Epstein and Zin (1989), developed for continuous time by Duffie and Epstein (1992). Utility can be represented by

$$V_t = E_t \int_t^\infty f(C_s, V_s) ds, \quad (\text{B.1})$$

where

$$f(C, V) = \frac{\beta}{1 - \frac{1}{\psi}} \frac{C^{1 - \frac{1}{\psi}} - ((1 - \gamma)V)^{\frac{1}{\theta}}}{((1 - \gamma)V)^{\frac{1}{\theta} - 1}}, \quad (\text{B.2})$$

and $\theta = (1 - \gamma)/(1 - \frac{1}{\psi})$. For $\psi = 1$, we assume

$$f(C, V) = \beta(1 - \gamma)V \left(\log C - \frac{1}{1 - \gamma} \log((1 - \gamma)V) \right). \quad (\text{B.3})$$

Here, $\beta > 0$ is the rate of time preference and $\gamma > 0$ is the coefficient of relative risk aversion.²⁹ The parameter ψ is commonly interpreted as the elasticity of intertemporal substitution and we will refer follow convention and refer to it as such. However, the nomenclature is misleading, since it will not in fact represent the sensitivity of consumption to changes in interest rates in the models we consider (though it play a principle role in controlling the sensitivity of consumption to changes in interest rates). When $1/\psi = \gamma$, the recursion is linear and these equations collapse to time-additive utility.

The expression for the state-price density in this economy is given by

$$\pi_t = \exp \left\{ \int_0^t \frac{\partial}{\partial V} f(C_s, V_s) ds \right\} \frac{\partial}{\partial C} f(C_t, V_t). \quad (\text{B.4})$$

²⁹Note that this form of the utility function does not assume $f(C, V)$ has the interpretation of period utility.

C Details for the multiperiod model (Section 3.7)

In this model, aggregate consumption growth is assumed to follow the process

$$\frac{dC_t}{C_t} = \mu_t dt + \sigma dB_t \quad (\text{C.1})$$

with

$$d\mu_t = \kappa_\mu(\bar{\mu} - \mu_t) dt + Z_t dN_t. \quad (\text{C.2})$$

To calculate the long-run impact of a shock Z_t , we assume that $N_t = 1$, and that there are no other shocks before some future time $t + \tau$. Then μ_s evolves according to

$$d\mu_s = \kappa_\mu(\bar{\mu} - \mu_s) ds \quad t < s \leq t + \tau.$$

Therefore $\mu_{t+\tau} = \bar{\mu} + Z_t e^{-\kappa_\mu \tau}$. It follows that the difference in expected consumption growth with and without the disaster equals:

$$\begin{aligned} E_t [\log C_{t+\tau} - \log C_t \mid N_s = 0, s < t + \tau] - (\bar{\mu} - \frac{1}{2}\sigma^2)\tau &= Z_t \int_0^\tau e^{-\kappa_\mu s} ds \\ &= \frac{Z_t}{\kappa_\mu} (1 - e^{-\kappa_\mu \tau}). \end{aligned}$$

Taking the limit as τ goes to infinity implies that Z_t/κ_μ is the long-run impact of a shock on consumption.

As in any dynamic model, returns are potentially sensitive to the timing of dividend payments. For this reason, we do not assume $D_t = C_t^\phi$ (which implies a mean of dividend growth that is significantly higher than in the data), but rather that dividends follow the process

$$\frac{dD_t}{D_t} = (\bar{\mu}_D + \phi(\mu_t - \bar{\mu})) dt + \phi\sigma dB_t.$$

The full solution for the price-dividend ratio (which is exact in the case of time-additive

utility and unit EIS and approximate otherwise), is given by

$$\frac{S_t}{D_t} = G(\mu_t) = \int_0^\infty e^{a_\phi(\tau) + b_\phi(\tau)\mu_t} d\tau$$

where

$$b_\phi(\tau) = \frac{\phi - \frac{1}{\psi}}{\kappa_\mu} (1 - e^{-\kappa_\mu \tau}) \quad (\text{C.3})$$

and

$$\begin{aligned} a_\phi(\tau) = & \left(\bar{\mu}_D - \frac{1}{\psi} \bar{\mu} - \beta + \frac{1}{2} \gamma \left(1 + \frac{1}{\psi} \right) \sigma^2 - \gamma \phi \sigma^2 \right. \\ & \left. + \lambda E_\nu \left[\left(\frac{1}{\theta} - 1 \right) \left(e^{(i_1 + \kappa_\mu)^{-1} (1 - \gamma) Z_t} - 1 \right) + \left(e^{(b_\phi(u) + (i_1 + \kappa_\mu)^{-1} (\frac{1}{\psi} - \gamma)) Z_t} - 1 \right) \right] \right) \tau - b_\phi(\tau) \bar{\mu} \end{aligned}$$

In this model, the response of equity to a disaster is equal to

$$\begin{aligned} Z_{S,t} &= \log \left(\frac{G(\mu_t + Z_t)}{G(\mu_t)} \right) \\ &= \log \left(\frac{1}{G(\mu_t)} \int_0^\infty e^{a_\phi(\tau) + b_\phi(\tau)(\mu_t + Z_t)} d\tau \right) \end{aligned} \quad (\text{C.4})$$

This expression seems complicated, but conceptually it is relatively simple. Equity value is a integral (sum) of discounted future dividends as in (24). The total decline is a weighted average of the decline in each of the terms in the integral. Equity prices fall on the onset of a disaster if and only if each term in the integral falls, which occurs if and only if $\phi > 1/\psi$.³⁰ That is, leverage must exceed the inverse of the EIS. For κ_μ relatively large, (C.3) is independent of τ and the equity response to a disaster can be approximated by

$$Z_{S,t} \approx \left(\phi - \frac{1}{\psi} \right) \frac{Z_t}{\kappa_\mu}. \quad (\text{C.5})$$

³⁰Note that equity prices fall if and only if

$$\int_0^\infty e^{a_\phi(\tau) + b_\phi(\tau)(\mu_t + Z_t)} d\tau < \int_0^\infty e^{a_\phi(\tau) + b_\phi(\tau)\mu_t} d\tau$$

Each term on the right hand side is lower than the corresponding term on the left hand side if and only if $b_\phi(\tau) > 0$, namely if and only if $\phi > 1/\psi$.

This result, combined with Proposition 1 implies the results in the main text.

D Details for the model with time-varying probability of disaster (Section 4.1)

In this model, the price-dividend ratio on the aggregate market satisfies

$$\frac{S_t}{D_t} = G(\lambda_t) = \int_0^\infty e^{a_\phi(\tau) + b_\phi(\tau)\lambda_t} d\tau.$$

These functions satisfy ordinary differential equations

$$a'_\phi(\tau) = \mu_D - \frac{1}{\psi}\mu - \beta + \left(\frac{1}{2}\left(1 + \frac{1}{\psi}\right) - \phi\right)\gamma\sigma^2 + \kappa\bar{\lambda}b_\phi(\tau) \quad (\text{D.1})$$

$$b'_\phi(\tau) = \frac{1}{2}\sigma_\lambda^2 b_\phi(\tau)^2 + \left(\left(1 - \frac{1}{\theta}\right)b\sigma_\lambda^2 - \kappa\right)b_\phi(\tau) + \frac{1}{2}\frac{1}{\theta}\left(\frac{1}{\theta} - 1\right)b^2\sigma_\lambda^2 + E_\nu \left[\left(\frac{1}{\theta} - 1\right) \left(e^{(1-\gamma)Z_t} - 1\right) + e^{(\phi-\gamma)Z_t} - 1 \right]. \quad (\text{D.2})$$

with boundary conditions $a_\phi(0) = b_\phi(0) = 0$. The system (D.1-D.2) has a closed-form solution (see Wachter (2013)).

Note that if the last two terms in this differential equation (D.2) were zero, then $b_\phi(\tau)$ would be identically zero, the price-dividend ratio would be a constant, and there would be no excess volatility or return predictability. We therefore seek to understand when the sum of these terms is negative, the empirically relevant case. We can recover these terms by evaluating the derivative of $b_\phi(\tau)$ at zero:³¹

$$b'_\phi(0) = \frac{1}{2}\frac{1}{\theta}\left(\frac{1}{\theta} - 1\right)b^2\sigma_\lambda^2 + E_\nu \left[\left(\frac{1}{\theta} - 1\right) \left(e^{(1-\gamma)Z_t} - 1\right) + e^{(\phi-\gamma)Z_t} - 1 \right] \quad (\text{D.3})$$

³¹The precise connection between the sign of $b'_\phi(0)$ and the dynamic behavior of the price-dividend ratio is discussed by Wachter (2013).

We now decompose the term inside the expectations as in Section 3.6:

$$\begin{aligned}
 b'_\phi(0) = & \underbrace{\frac{1}{2} \frac{1}{\theta} \left(\frac{1}{\theta} - 1 \right) b^2 \sigma_\lambda^2 - E_\nu \left[\left(1 - \frac{1}{\theta} \right) (e^{(1-\gamma)Z_t} - 1) - (e^{-\gamma Z_t} - 1) \right]}_{\text{risk-free rate}} \\
 & - \underbrace{E_\nu [(e^{-\gamma Z_t} - 1)(1 - e^{\phi Z_t})]}_{\text{equity premium}} + \underbrace{E_\nu [e^{\phi Z_t} - 1]}_{\text{expected dividend growth}}. \quad (\text{D.4})
 \end{aligned}$$

Equation (D.4) shows that for time-additive utility and for $\psi = 1$, the trade-offs determining whether prices rise or fall in response to changes in the disaster probability are exactly the same as the comparative statics result in Section 3.6 (for $\psi \neq 1$ and $\psi \neq 1/\gamma$, there is a slight difference arising from the additional term in the riskfree rate expression).

This closed-form analysis illustrates the features of a calibration that are likely to explain stock market data. First, it almost goes without saying that some risk aversion will be necessary to achieve an equity premium. In order to generate sufficient stock market volatility along with a low and stable riskfree rate, it is helpful to have dividends that are more responsive to disasters than consumption ($\phi > 1$), a relatively high EIS ($\psi \geq 1$), and, separately, a preference for early resolution of uncertainty ($\gamma > 1/\psi$).

References

- Abel, Andrew, 1999, Risk premia and term premia in general equilibrium, *Journal of Monetary Economics* 43, 3–33.
- Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters Implied by Equity Index Options, *The Journal of Finance* 66, 1969–2012.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long-run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866.
- Barro, Robert J., 2009, Rare disasters, asset prices, and welfare costs, *American Economic Review* 99, 243–264.
- Barro, Robert J., and José F. Ursúa, 2008, Macroeconomic crises since 1870, *Brookings Papers on Economic Activity* no. 1, 255–350.
- Beeler, Jason, and John Y. Campbell, 2012, The long-run risks model and aggregate asset prices: An empirical assessment, *Critical Finance Review* 1, 141–182.
- Berkman, Henk, Ben Jacobsen, and John B. Lee, 2011, Time-varying rare disaster risk and stock returns, *Journal of Financial Economics* 101, 313–332.
- Bhamra, Harjoat S., and Ilya A. Strebulaev, 2011, The effects of rare economic crises on credit spreads and leverage, Working paper.
- Boldrin, Michele, Lawrence J. Christiano, and Jonas DM Fisher, 2001, Habit persistence, asset returns, and the business cycle, *American Economic Review* pp. 149–166.
- Bollerslev, Tim, and Viktor Todorov, 2011, Tails, Fears, and Risk Premia, *The Journal of Finance* 66, 2165–2211.

- Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael R. Roberts, 2007, On the importance of measuring payout yield: Implications for empirical asset pricing, *Journal of Finance* 62, 877–915.
- Breenden, Douglas T., 1979, An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265–296.
- Brown, Stephen J., William N. Goetzmann, and Stephen A. Ross, 1995, Survival, *The Journal of Finance* 50, 853–873.
- Calvet, Laurent E., and Adlai J. Fisher, 2007, Multifrequency news and stock returns, *Journal of Financial Economics* 86, 178–212.
- Campbell, John Y., 2003, Consumption-based asset pricing, in G. Constantinides, M. Harris, and R. Stulz, eds.: *Handbook of the Economics of Finance, vol. 1b* (Elsevier Science, North-Holland).
- Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- Campbell, John Y., and Robert J. Shiller, 1991, Yield spreads and interest rate movements: A bird’s eye view, *Review of Economic Studies* 58, 495–514.
- Chabi-Yo, Fousseni, Stefan Ruenzi, and Florian Weigert, 2014, Crash sensitivity and the cross-section of expected stock returns, Working paper.
- Chen, Hui, Scott Joslin, and Ngoc-Khanh Tran, 2012, Rare disasters and risk sharing with heterogeneous beliefs, *Review of Financial Studies* 25, 2189–2224.

- Cochrane, John H., 1992, Explaining the variance of price-dividend ratios, *Review of Financial Studies* 5, 243–280.
- Cochrane, John H., and Monika Piazzesi, 2005, Bond risk premia, *American Economic Review* 95, 138–160.
- Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer, 2013, Parameter Learning in General Equilibrium: The Asset Pricing Implications, National Bureau of Economic Research Working Paper 19075.
- Constantinides, George M., 2008, Discussion of “Macroeconomic crises since 1870”, *Brookings Papers on Economic Activity* no. 1, 336–350.
- Diamond, Peter A., 1982, Wage determination and efficiency in search equilibrium, *The Review of Economic Studies* 49, 217–227.
- Du, Du, 2011, General equilibrium pricing of options with habit formation and event risks, *Journal of Financial Economics* 99, 400–426.
- Duffee, Gregory R., 2013, Bond pricing and the macroeconomy, in George M. Constantinides, Milton Harris, and Rene M. Stulz, eds.: *Handbook of the Economics of Finance*, vol. 2b (Elsevier, North-Holland).
- Duffie, Darrell, 2001, *Dynamic Asset Pricing Theory*. (Princeton University Press Princeton, NJ) 3 edn.
- Duffie, Darrell, and Larry G. Epstein, 1992, Stochastic differential utility, *Econometrica* 60, 353–394.
- Epstein, Larry, and Stan Zin, 1989, Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.
- Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.

- Farhi, Emmanuel, and Xavier Gabaix, 2014, Rare Disasters and Exchange Rates, Working paper.
- Gabaix, Xavier, 2008, Linearity-generating processes: A modelling tool yielding closed forms for asset prices, Working paper, New York University.
- Gabaix, Xavier, 2011, Disasterization: A simple way to fix the asset pricing properties of macroeconomic models, *American Economic Review* 101, 406–409.
- Gabaix, Xavier, 2012, An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.
- Gao, George P., and Zhaogang Song, 2013, Rare disaster concerns everywhere, Working paper, Cornell University.
- Gao, George P., Zhaogang Song, and Liyan Yang, 2014, Perceived crash risk and cross-sectional stock returns, Working paper.
- Gillman, Max, Michal Kejak, and Michal Pakoš, 2014, Learning about rare disasters: Implications for consumption and asset prices, forthcoming, *Review of Finance*.
- Goetzmann, William N., and Philippe Jorion, 1999, Re-emerging markets, *Journal of Financial and Quantitative Analysis* pp. 1–32.
- Gourio, François, 2009, Disasters risk and business cycles, National Bureau of Economic Research, Working paper 15399.
- Gourio, François, 2012, Disaster risk and business cycles, *American Economic Review* 102, 2734–2766.
- Gourio, François, 2013, Credit risk and disaster risk, *American Economic Journal: Macroeconomics* 5, 1–34.
- Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 221–273.


- Hansen, Lars Peter, 2007, Beliefs, doubts and learning: Valuing macroeconomic risk, *The American Economic Review* 97, 1–30.
- Harrison, Michael, and David Kreps, 1979, Martingales and multiperiod securities markets, *Journal of Economic Theory* 20, 381–408.
- Jermann, Urban J., 1998, Asset pricing in production economies, *Journal of Monetary Economics* 41, 257–275.
- Julliard, Christian, and Anisha Ghosh, 2012, Can rare events explain the equity premium puzzle?, *The Review of Financial Studies* 25, 3037–3076.
- Kaltenbrunner, G., and L. A. Lochstoer, 2010, Long-run risk through consumption smoothing, *Review of Financial Studies* 23, 3190–3224.
- Keim, Donald B., and Robert F. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357–390.
- Kelly, Bryan, and Hao Jiang, 2014, Tail risk and asset prices, forthcoming, *Review of Financial Studies*.
- Larrain, Borja, and Motohiro Yogo, 2008, Does firm value move too much to be justified by subsequent changes in cash flow?, *Journal of Financial Economics* 87, 200–226.
- LeRoy, Stephen F., and Richard D. Porter, 1981, The present-value relation: Tests based on implied variance bounds, *Econometrica* 49, 555–574.
- Lettau, Martin, and Harald Uhlig, 2000, Can habit formation be reconciled with business cycle facts?, *Review of Economic Dynamics* 3, 79–99.
- Lettau, Martin, and Jessica A. Wachter, 2011, The term structures of equity and interest rates, *Journal of Financial Economics* 101, 90 – 113.
- Longstaff, Francis A., and Monika Piazzesi, 2004, Corporate earnings and the equity premium, *Journal of Financial Economics* 74, 401–421.

- Maddison, Angus, 2003, *The World Economy: Historical Statistics*. (OECD Paris).
- Malmendier, Ulrike, and Stefan Nagel, 2011, Depression babies: Do macroeconomic experiences affect risk taking?*, *Quarterly Journal of Economics* 126, 373–416.
- Manela, Asaf, and Alan Moreira, 2013, News implied volatility and disaster concerns, Working paper, Washington University and Yale University.
- Martin, Ian, 2013, The Lucas Orchard, *Econometrica* 81, 55–111.
- Mehra, Rajnish, and Edward Prescott, 1985, The equity premium puzzle, *Journal of Monetary Economics* 15, 145–161.
- Mehra, Rajnish, and Edward Prescott, 1988, The equity risk premium: A solution?, *Journal of Monetary Economics* 22, 133–136.
- Merton, Robert C., 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.
- Mortensen, Dale T, and Christopher A Pissarides, 1994, Job creation and job destruction in the theory of unemployment, *The Review of Economic Studies* 61, 397–415.
- Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa, 2013, Crises and recoveries in an empirical model of consumption disasters, *American Economic Journal: Macroeconomics* 5, 35–74.
- Nowotny, Michael, 2011, Disaster begets crisis: The role of contagion in financial markets, Working paper, Boston University.
- Petrosky-Nadeau, Nicolas, Lu Zhang, and Lars-Alexander Kuehn, 2013, Endogenous economic disasters and asset prices, Fisher College of Business Working Paper.
- Rietz, Thomas A., 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117–131.

- Seo, Sang Byung, 2014, Correlated defaults and economic catastrophes: Link between CDS market and asset returns, working paper.
- Seo, Sang Byung, and Jessica A. Wachter, 2015, Option prices in a model with stochastic disaster risk, NBER Working Paper #19611.
- Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421–436.
- Stambaugh, Robert F., 1988, The information in forward rates: Implications for models of the term structure, *Journal of Financial Economics* 21, 41–70.
- Tsai, Jerry, 2014, Rare disasters and the term structure of interest rates, working paper.
- Tsai, Jerry, and Jessica A. Wachter, 2014a, Pricing long-lived securities in dynamic economies, working paper, University of Oxford and University of Pennsylvania.
- Tsai, Jerry, and Jessica A. Wachter, 2014b, Rare booms and disasters in a multi-sector endowment economy, NBER Working Paper No. 20062.
- Verdelhan, Adrien, 2015, Empirical Evidence on Exchange Rate Movements, *Annual Review of Financial Economics* 7.
- Veronesi, Pietro, 2004, The Peso problem hypothesis and stock market returns, *Journal of Economic Dynamics and Control* 28, 707–725.
- Wachter, Jessica A., 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *The Journal of Finance* 68, 987–1035.
- Weil, Philippe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 402–421.
- Weil, Philippe, 1990, Nonexpected utility in macroeconomics, *Quarterly Journal of Economics* 105, 29–42.

Weitzman, Martin L., 2007, Subjective expectations and asset-return puzzles, *American Economic Review* 97, 1102–1130.

Figure 1: Distribution of consumption declines in the event of a disaster



Notes: Histograms show the distribution of large consumption declines ($1 - e^Z$) in percentages. Panel A shows data for 24 countries, 17 of which are OECD countries and seven of which are not; Panel B shows data for the subsample of OECD countries. Data are from Barro and Ursua (2008). The cut-off for the definition of a disaster is 15%.

Table 1: Rates of return for the i.i.d. model

	No disaster	No default	Default
Panel A: Treasury Bill return			
Time-additive utility ($\psi = 1/4$)	10.64	-3.36	-0.25
$\psi = 1$	4.89	-2.87	0.23
$\psi = 2$	3.93	-2.79	0.31
Panel B: Equity premium, population			
$\psi = 1/4, 1, 2$	0.06	7.22	4.36
Panel C: Equity premium, observed			
$\psi = 1/4, 1, 2$	0.06	7.82	4.72

Notes: The table shows the rate of return on the government bill (Panel A), the population equity premium (Panel B) and the equity premium that would be observed in a sample without disasters (Panel C). Parameters (in annual terms) are as follows: disaster probability $\lambda = 0.0218$; discount rate $\beta = 0.03$; risk aversion $\gamma = 4$; normal-times consumption growth $\mu = 0.0195$, and volatility $\sigma = 0.0125$; leverage $\phi = 1$; and (when applicable) default probability conditional on disaster $q = 0.4$. For T-Bill returns in the case with default, we report the yield (equivalently, the return if default is not realized). All the results are in annual percentage terms.

Table 2: The equity premium in the multiperiod disaster model

	Consumption claim ($\phi = 1$)		Dividend claim ($\phi = 3$)	
	Population EP	Observed EP	Population EP	Observed EP
Time-additive utility ($\psi = 1/4$)	0.06	-6.13	0.19	-0.73
$\psi = 1$	0.06	0.06	4.19	5.15
$\psi = 2$	2.69	3.01	6.94	8.05

Notes: The equity premium for the consumption and dividend claim in the multiperiod disaster model. Column 2 and 3 report the population equity premium and the equity premium that would be observed in a sample without disasters for the consumption claim ($\phi = 1$), and Column 4 and 5 report premiums for the dividend claim ($\phi = 3$). Parameters are as follows: disaster probability $\lambda = 0.0218$, discount rate $\beta = 0.01$, risk aversion $\gamma = 4$, normal time consumption and dividend growth $\mu = \mu_D = 0.0195$, volatility $\sigma = 0.0125$, mean-reversion parameter in the μ_t process $\kappa_\mu = 1$. All the results are in annual percentage terms.

Table 3: Return moments in the time-varying disaster probability model

	Time-varying disaster probability					Constant disaster probability			
	Data	No-disaster simulations				No-disaster simulations			
		0.05	0.50	0.95	Pop.	0.05	0.50	0.95	Pop.
Panel A: Time-additive utility ($\psi = 1/3$)									
$E[R^b]$	1.24	0.08	3.55	5.26	1.79	1.93	1.93	1.93	1.69
$\sigma(R^b)$	2.57	1.51	3.03	6.00	6.00	0.00	0.00	0.00	2.89
$E[R^e - R^b]$	7.69	1.51	3.79	7.40	4.15	3.98	5.35	6.73	4.25
$\sigma(R^e)$	17.72	5.71	6.68	7.69	11.58	5.72	6.68	7.68	11.65
Sharpe ratio	0.44	0.23	0.57	1.13	0.36	0.59	0.80	1.04	0.36
Panel B: $\psi = 1$									
$E[R^b]$	1.24	-0.04	1.48	2.22	0.58	0.81	0.81	0.81	0.56
$\sigma(R^b)$	2.57	0.65	1.31	2.65	3.78	0.00	0.00	0.00	2.90
$E[R^e - R^b]$	7.69	4.74	7.04	10.65	7.17	3.93	5.28	6.63	4.13
$\sigma(R^e)$	17.72	10.75	15.05	21.34	19.73	5.65	6.60	7.61	11.58
Sharpe ratio	0.44	0.35	0.47	0.61	0.36	0.59	0.80	1.04	0.36
Panel C: $\psi = 2$									
$E[R^b]$	1.24	-0.80	0.61	1.29	-0.27	0.53	0.53	0.53	0.28
$\sigma(R^b)$	2.57	0.60	1.20	2.45	3.68	0.00	0.00	0.00	2.88
$E[R^e - R^b]$	7.69	5.50	7.91	11.83	8.19	3.91	5.26	6.64	4.16
$\sigma(R^e)$	17.72	12.04	17.13	24.60	22.07	5.64	6.58	7.57	11.48
Sharpe ratio	0.44	0.36	0.47	0.59	0.37	0.59	0.80	1.04	0.36

Notes: Parameters are as follows: average disaster probability $\bar{\lambda} = 0.0218$, discount rate $\beta = 0.01$, risk aversion $\gamma = 3$, normal-times consumption growth $\mu = 0.0195$, consumption growth volatility $\sigma = 0.0125$, dividend growth $\mu_D = 0.04$, idiosyncratic volatility $\sigma_i = 0.05$, leverage $\phi = 3$, and mean-reversion $\kappa_\lambda = 0.12$. Volatility $\sigma_\lambda = 0.08$ in the time-varying disaster probability case, and $\sigma_\lambda = 0$ in the constant disaster probability case.