

The Probability of Rare Disasters: Estimation and Implications

Emil Siriwardane

Working Paper 16-061



The Probability of Rare Disasters: Estimation and Implications

Emil Siriwardane
Harvard Business School

Working Paper 16-061

Copyright © 2015 by Emil Siriwardane

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

The Probability of Rare Disasters: Estimation and Implications*

Emil Siriwardane[†]

October 2015

Abstract

I analyze a rare disasters economy that yields a measure of the risk neutral probability of a macroeconomic disaster, p_t^* . A large panel of options data provides strong evidence that p_t^* is the single factor driving option-implied jump risk measures in the cross section of firms. This is a core assumption of the rare disasters paradigm. A number of empirical patterns further support the interpretation of p_t^* as the risk-neutral likelihood of a disaster. First, standard forecasting regressions reveal that increases in p_t^* lead to economic downturns. Second, disaster risk is priced in the cross section of U.S. equity returns. A zero-cost equity portfolio with exposure to disasters earns risk-adjusted returns of 7.6% per year. Finally, a calibrated version of the model reasonably matches the: (i) sensitivity of the aggregate stock market to changes in the likelihood of a disaster and (ii) loss rates of disaster risky stocks during the 2008 financial crisis.

*I am especially thankful to Xavier Gabaix for his advice on this research. In addition, Viral Acharya, David Backus, Malcolm Baker, Robert Barro, Juliane Begenau, John Campbell, Lauren Cohen, Robert Engle, Sam Hanson, Nikunj Kapadia, Matteo Maggiori, Chris Malloy, Alexi Savov, Erik Stafford, Adi Sunderam, and Stijn Van Nieuwerburgh provided many valuable comments. I also thank seminar participants at Harvard Economics (Macro), NYU Stern, the MFM Spring 2013 Meetings, and the Office of Financial Research. Finally, I am grateful for grant support from the MFM and the Center for Global Economy and Business at NYU Stern.

[†]Affiliation: Harvard Business School. Email: esiriwardane@hbs.edu

1 Introduction

A longstanding puzzle in macroeconomics and financial economics is that consumption growth is not volatile enough to justify why the aggregate stock market earns large returns in excess of riskless bonds. In response, [Rietz \(1988\)](#) and [Barro \(2006\)](#) argue that the economy is subject to rare “disasters” such as wars and depressions; in turn, the risk premium on stocks is not surprising because stocks are highly exposed to these rare events.

While this argument is simple, the empirical assessment of these models has proved more controversial. [Barro and Ursua \(2008\)](#) use personal consumption data to estimate the average likelihood and severity of a crisis, finding evidence that supports the original [Rietz \(1988\)](#) and [Barro \(2006\)](#) claim. [Nakamura et al. \(2013\)](#) show that rare disasters can explain stock market performance even if disasters unfold over a number of years and are followed by economic recoveries. On the other hand, many argue that rare disaster models are not supported by equity or options data ([Julliard and Ghosh \(2012\)](#), [Backus, Chernov, and Martin \(2011\)](#)), and that more generally, the probability of a disaster may not even be identifiable ([Chen et al. \(2015\)](#)).

A critical assumption made by all of the aforementioned research is that the frequency of disasters is time-invariant. However, [Gabaix \(2012\)](#) and [Seo and Wachter \(2015\)](#) show theoretically that much of the empirical debate can be resolved by allowing the likelihood and/or severity of a disaster to be time varying. [Gabaix \(2012\)](#) also demonstrates that variable disasters can explain a host of other asset pricing patterns like the volatility of price-dividend ratios and return predictability. Nonetheless, there is not a great deal of evidence on whether disasters are indeed time-varying, as this variation is notoriously difficult to estimate. Moreover, the ability to detect time-series variation in the likelihood of a disaster is important because this yields additional ways to evaluate the models outside of matching unconditional average returns.

In this paper, I develop a novel approach to measure time-series variation in the Arrow-

Debreu state price of a disaster, or the so-called “risk-neutral” probability of a disaster. Instead of using aggregate stock price or macroeconomic data as in much of the previous work, my strategy takes advantage of information contained in the cross section of option and equity returns. I focus on the risk-neutral probability of a disaster because option data contains information on the risk-neutral process that cannot be translated to realized likelihoods without making stronger assumptions on preferences.¹ This is not particularly limiting given that it is the risk-neutral likelihood of a disaster that dictates future economic outcomes (e.g. investment decisions or asset returns).

The high-level logic of my approach is as follows. The core assumption of rare disaster models is that all assets in the economy are exposed to an aggregate jump-risk factor. If this is the case, then time-series variation in the tails of individual asset returns should be driven by this single factor, even though cross-sectionally some assets are more exposed to the tail event than others. I formalize this logic in the context of a standard rare disasters model, which delivers a straightforward way to use stock option prices to diagnose the factor structure of asset tails. Options are a natural asset class for this problem because they are a forward-looking measure of firm performance and do not require realizations of disasters to occur. Another advantage of my approach is that I use a non-parametric measure of jump risk in options. This lets me test for the presence of a time-varying disaster process without imposing any a priori structure on the data.

Consistent with the basic premise of the model, a single factor captures 50% of the time-series variation in a large panel of option-implied jump measures. Using the model, I interpret this factor as measuring temporal variation in the risk-neutral probability of a disaster, p_t^* . Qualitatively, p_t^* evolves as expected, peaking during tumultuous times like the 2003 Iraq War

¹i.e. I estimate time-series variation of the (scaled) Arrow-Debreu security that only pays off in the disaster state, not the actual likelihood of this state.

and the 2008 U.S. financial crisis. As posited in the literature, p_t^* is clearly time-varying and persistent, with an AR(1) coefficient of nearly 0.91 for monthly data. Importantly, I do not pin down the mean of p_t^* , but its time-series variation is still valuable because rare disaster models contain many predictions based only on the time-series properties of disaster risk.

For instance, [Gourio \(2012\)](#) and [Kilic and Wachter \(2015\)](#) build general equilibrium models where an exogenous increase in the likelihood of a disaster leads to endogenous declines in macroeconomic activity, even without the realization of a disaster. Consistent with these models, from April 1996 to April 2015, an increase in p_t^* forecasts significant reductions in employment, industrial production, total capacity utilization, and manufacturing activity. Moreover, the predictive power of p_t^* outperforms standard financial market indicators like the shape of the Treasury yield curve, the real federal funds rate, and aggregate stock market performance.

The rare disasters paradigm also makes a strong prediction for how time-series variation in the probability of disaster relates to the cross-sectional performance of stocks. Intuitively, stocks that rise in value as disasters become more likely are valuable hedges against macroeconomic crises. As a result, these stocks earn low risk premiums relative to stocks that fall in value as disasters become more imminent. Still, the operational assumption behind this intuition is that one can empirically measure the comovement of stocks with the probability of a disaster. This is another reason why my estimate of p_t^* is so useful.

To test the model's prediction that disaster risky stocks require a premium, I implement a simple real-time trading strategy. Specifically, I sort stocks into portfolios based on how much they comove with changes in p_t^* . As implied by the model, I find that disaster risky stocks — those that lose value as p_t^* increases — outperform low disaster risk stocks by nearly 7.6% per year, even after accounting for exposure to standard risk factors. In other words, disaster risk appears to be priced in the cross section of U.S. equities, and the premium for bearing disaster

risk is about 7.6% annually.

In the final part of the paper, I make stronger assumptions about the mean frequency and severity of disasters, as well as preferences. This allows me to quantitatively compare a calibrated version of the model to aggregate stock market movements and the cross section of equity returns. For example, I ask the model how much the stock market should move in response to changes in p_t^* . Empirically, I find that when p_t^* increases, the market declines by an amount that is quantitatively matched by the calibrated model. This finding is also in line with [Barro \(2009\)](#), who shows that disaster models with Epstein-Zin-Weil preferences require an elasticity of substitution greater than one in order for the aggregate stock market to decline as disasters become more likely.

In a similar fashion, the calibrated model also yields predictions for how much disaster risky stocks must fall in a crisis in order to justify their large returns. I then compare this to how much these stocks actually fall during the 2008 U.S. financial crisis. Along this dimension, the data also matches the model reasonably well. In the model, disaster risky stocks must fall by 57% in a disaster to warrant their high average returns. During the crisis, these same stocks lose about 51% of their value, though many of these stocks benefited from the unprecedented intervention of the government into U.S. capital markets.

Literature Review

The first major contribution of this paper is to estimate time-series variation in the risk-neutral probability of a disaster, and the primary asset class that I use is the options market. As mentioned earlier, much of the existing research begins from the assumption that the disaster arrival rate is i.i.d. [Backus, Chernov, and Martin \(2011\)](#) use option prices on the S&P 500 to estimate a structural model under this i.i.d assumption. Their findings show that options markets

imply that disasters happen much more frequently (yet with less severity) than suggested by consumption data (Barro (2006) or Barro and Ursua (2008)). However, Seo and Wachter (2015) note that with persistent and time-varying disasters, parameters from consumption data on disasters can actually generate realistic option prices.² Thus, drawing conclusions about rare disasters from option prices depends heavily on model specification, particularly on the time-varying nature of disasters.

To circumvent any misspecification issues, my estimate of time-series variation in the probability of disaster is non-parametric and uses options on a large cross section of firms; at the firm level, my jump risk measure applies to very general return processes and is not necessarily linked to a disaster model. This is useful because it allows me to first characterize the option-implied tails of many different firms, and then check whether the data is consistent with the disaster model. In turn, the fact that this cross section of jump-risk measures is driven by a common and time-varying factor is in strong agreement with, among others, Barro (2006), Gabaix (2012), and Seo and Wachter (2015).

The second major contribution of this paper is to provide evidence in favor of the rare disasters paradigm using the cross section of U.S. equity returns. Much of the literature has focused on whether disasters can rationalize the returns of the aggregate stock market, though it is important to ask whether these models can also explain the cross section of stock returns. Gao and Song (2015) find evidence that jump risk is priced across a wide range of asset classes, though their focus is on broad market indices instead individual stocks. My finding that disaster risk is priced in the cross section of U.S. equities is also consistent with Cremers, Halling, and Wein-

²The intuition of their findings is as follows: when disasters are time-varying, movements in the probability of the disaster lead to movements in price-dividend ratios. This generates volatile price-dividend ratios, even if the disaster never occurs. Without this mechanism, the volatility of price-dividend ratios in normal times appears too low relative to the volatility of consumption growth in normal times. In turn, models without time-varying disaster risk will generate large pricing errors for options with strikes that are near the current spot price (at-the-money).

baum (2015), however their focus is on high-frequency (i.e. daily) jumps, not macroeconomic disasters.³ Perhaps the closest companion to this paper is Farhi, Fraiburger, Gabaix, Ranciere, and Verdelhan (2014), who study how disasters impact foreign exchange markets. These authors use a version of the rare disaster model to link currency option premiums to carry trade premiums. Importantly, I too map my empirical findings to a standard rare disasters model. This provides me an additional way to evaluate the rare disaster model because I can compare the cross-sectional pricing results with the predictions of a calibrated version of the model. Along this dimension, the rare disaster model also holds up reasonably well to the data.

2 Theoretical Motivation

In this section, I first present a non-parametric measure of jump risk that can be constructed using observed option prices on any given firm. I then use a standard model of time-varying disasters to map this jump-risk measure to model quantities. Through this process, it is easy to see that rare disaster models impose a tight restriction on how jump-risk measures should comove across firms. This restriction then forms the basis of my empirical work in Section 3.

2.1 A Non-Parametric Measure of Jump Risk

As mentioned, my estimate of time-series variation in the probability of a disaster is non-parametric. To develop this idea further, I need to first introduce the concept of the risk-neutral probability measure and the physical probability measure. The physical probability measure describes different states of the world according to their actual likelihood of occurring. The risk-neutral

³Other evidence of tail risk being priced in the cross section of equities can be found in Kelly and Jiang (2014), who rely on extreme value theory to learn about the tails of asset prices from realized stock returns (not options). Bollerslev, Todorov, and Xu (2015) also use options to relate jump risk to stock returns. They use S&P 500 options to parametrize option-implied tails and find that their estimated jump risk premia forecast future stock returns.

measure instead weights states of nature according to marginal utility in those states. Naturally, states in which marginal utility is high (i.e. recession states) have a higher risk-neutral likelihood than a physical likelihood. In fact, risk-neutral probabilities are just Arrow-Debreu state prices, appropriately normalized so that they correspond to a probability measure. In any equilibrium model, the mapping between risk-neutral and actual probability measures is therefore determined entirely by the pricing kernel implied by the model.

Without imposing a specific asset pricing model, suppose that equilibrium stock returns for firm i evolve according to the following stochastic differential equation under the risk-neutral measure:

$$\frac{dP_{it}^*}{P_{it}^*} = \mu_{it}^* dt + \sigma_{it}^* dW_{it}^* + (F_{it}^* - 1) dJ_{it}^* \quad (1)$$

where P_{it} is the stock price of firm i at time t . Here, I am using the * superscript to make clear that we are examining returns under the risk-neutral probability measure. According to this reduced-form specification, risk-neutral returns are subject to two-shocks. The first is a diffusive Brownian motion shock, dW_{it}^* , where the risk-neutral volatility σ_{it}^* of the total diffusive portion of returns is left unspecified.

The second is a “jump” shock, captured by the process dJ_{it}^* . dJ_{it}^* takes a value of 1 if a jump occurs over the next instant, and 0 if not. The risk-neutral Poisson arrival rate of jumps is time varying, firm specific, and denoted by p_{it}^* . Hence, a jump happens over a small time interval, dt , with a risk-neutral probability of $p_{it}^* dt$. When a jump does occur, the size of the jump is governed by the stochastic variable F_{it}^* . For instance, if $F_{it}^* = 0.9$, then the stock price will drop by 10% when a jump occurs. In addition, I assume that F_{it}^* and p_{it}^* are independent, and that the risk-neutral distribution of F_{it}^* is time-invariant. Later, I discuss the importance of these assumptions when interpreting my empirical findings.

The risk-neutral return process in Equation (1) is intentionally quite general, and therefore

nesses the equilibrium outcome of many different asset pricing models.⁴ Importantly, different asset pricing models will lead to different restrictions on the shock processes, dW_{it}^* and dJ_{it}^* . I am specifically interested in what rare disaster models imply about the process dJ_{it}^* , both theoretically and empirically.

In this general case though, recent developments in option pricing theory suggest a way to use *observed* option prices to learn something about p_{it}^* . In particular, [Du and Kapadia \(2013\)](#) develop a measure of jump risk that can be inferred from option prices, which I denote by $\mathbb{D}_{it}(\tau)$ for “disaster risk”. $\mathbb{D}_{it}(\tau)$ is the value of a simple portfolio of call and put options on firm i and formed on date t . It consists of options written across many different strike prices, all with maturity of τ . The exact details of this measure can be found in [Appendix C](#).

To develop some intuition now, consider the popular VIX index. The VIX is the price of a portfolio of call and put options written on the S&P 500. In the *absence* of jumps, the value of the VIX (squared) is the expected risk-neutral variance of the S&P 500 over the next month, or $\Omega_{it}(\tau) := \mathbb{E}_t^* \left[\int_t^{t+\tau} (\sigma_{is}^*)^2 ds \right]$ for $i = \text{S\&P 500}$ and $\tau = 30$ days. However, when risk-neutral returns are subject to jumps, the VIX is a biased estimate of $\Omega_{it}(\tau)$.⁵ The reason is that jumps create additional return volatility, as squared returns can come from continuous diffusive movements or discontinuous jumps. As pointed out by [Du and Kapadia \(2013\)](#), it is possible to design an alternative option portfolio — one whose value is $\mathbb{D}_{it}(\tau)$ — that isolates the bias in the VIX that is induced by jumps. In turn, $\mathbb{D}_{it}(\tau)$ captures the “extra” variation that jumps create for risk-neutral returns. In [Appendix C](#), I also use a stylized example to deepen the intuition of why \mathbb{D}_{it} is able to isolate jumps.

More specifically, [Carr and Wu \(2009\)](#) and [Du and Kapadia \(2013\)](#) show that when risk-

⁴The only restriction is the standard no-arbitrage requirement that the discounted risk-neutral stock price is a martingale. This implies that $\mu_{it} = r_{ft} - p_{it}^* \mathbb{E}_t[F_{it} - 1]$, where r_{ft} is the instantaneous riskless rate.

⁵This point has been made by, among many others, [Martin \(2013\)](#) and [Carr and Wu \(2009\)](#).

neutral returns evolve according to Equation (1), $\mathbb{D}_{it}(\tau)$ captures the following notion of firm i 's jump risk from time t to $t + \tau$:

$$\mathbb{D}_{it}(\tau) = \Psi_i \times p_{it}^*(\tau)$$

With some abuse of notation, I use $p_{it}^*(\tau) := \mathbb{E}_t^* \left[\int_t^{t+\tau} p_{is}^* ds \right]$ to denote the risk-neutral expected likelihood of a jump from t to $t + \tau$. Ψ_i is a firm-specific constant that depends on the distribution of F_{it} . The exact expression for Ψ_i is given in Equation (47) of Appendix C.6.⁶

In my empirical work, I will use observed option prices to form a $\mathbb{D}_{it}(\tau)$ measure for a very large panel of firms. It should be clear by now that this measure is non-parametric in the sense that its construction and interpretation are not anchored to any specific asset pricing model. My next task is to map $\mathbb{D}_{it}(\tau)$ to a standard model of time-varying disasters. This will generate a key testable prediction for how $\mathbb{D}_{it}(\tau)$ measures should comove across firms.

2.2 Mapping to a Rare Disasters Model

I now solve a continuous time generalization of the endowment economy in [Gabaix \(2012\)](#). In the main text, I present only the basic ingredients of the model and the important equilibrium outcomes. The full details of the model can be found in Appendix A. The purpose of this exercise is to develop intuition for how the $\mathbb{D}_{it}(\tau)$ measures from the previous section should behave under the null of the model.

⁶To be clear, \mathbb{D}_{it} captures any expected *discontinuities* in the risk-neutral return process, though it is not possible to infer the size of the discontinuity (jump) from \mathbb{D}_{it} alone. An alternative notion of disaster risk could be to look at risk-neutral probabilities implied by deep out-of-the-money put options. The advantage of this approach is that one could say something about the size of the expected disaster, yet the disadvantage is that it requires one to observe deep enough out-of-the-money puts. In rare disaster models, the extreme negative tail risk of the economy is almost always modeled as a pure jump, in which case the two measures should coincide.

2.2.1 Model Primitives

Aggregate Consumption To start, the aggregate endowment evolves as follows:

$$\frac{dC_t}{C_t} = g_C dt + \sigma_C dW_{Ct} + (B_t - 1)dJ_t$$

where W_{Ct} is a standard Brownian motion and J_t is a jump process that represents macroeconomic disasters. dJ_t takes a value of one if a disaster occurs and zero otherwise. The Poisson arrival rate of disasters is time varying and denoted by p_t . Hence, a disaster happens over a small time interval, dt , with a probability $p_t dt$. The variable $B_t > 0$ captures the magnitude of a consumption disaster when it occurs. For example, a value of $B_t = 0.8$ means that when a disaster occurs, aggregate consumption drops by 20%. I further assume that the intensity process describing disaster arrivals, p_t , is independent of the magnitude of a disaster, B_t . The Brownian motion W_{Ct} is independent of J_t and B_t , which means that in the absence of disasters (i.e. normal times) consumption growth is i.i.d log-normal with a volatility of σ_C .

Preferences I assume the representative agent has recursive preferences as in [Epstein and Zin \(1989\)](#) and [Duffie and Epstein \(1992\)](#). The main advantage of recursive preferences is the ability to decouple the agent's elasticity of intertemporal substitution, ψ , from risk aversion, γ . In a discrete time setting, [Epstein and Zin \(1989\)](#) show that the representative agent's stochastic discount factor can be expressed as:

$$\frac{M_{t+\tau}}{M_t} = e^{-\frac{\rho}{\chi}\tau} \left(\frac{C_{t+\tau}}{C_t} \right)^{-\frac{1}{\chi\psi}} R_{C,t+\tau}^{1/\chi-1}$$

where $\chi = (1 - 1/\psi)/(1 - \gamma)$ and $R_{C,t+\tau}$ is the gross return of a claim to aggregate consumption between t and $t + \tau$. Here, ρ captures the rate of time preference of the agent. When solving the

model, I obtain the stochastic discount factor by taking the continuous time limit of the discrete time stochastic discount factor, i.e. as $\tau \searrow 0$.

Setup for Individual Stocks A generic stock i is defined as a claim to a stream of dividends $(D_{it})_{t \geq 0}$, where dividends grow according to:

$$\frac{dD_{it}}{D_{it^-}} = g_{iD}dt + dN_{it}^D + (F_{it} - 1)dJ_t$$

The dividend growth shock N_{it}^D is a martingale so that $\mathbb{E}_t[dN_{it}^D] = 0$ and is independent of the disaster event. The higher moments of N_{it}^D are left unspecified for now. In equilibrium, only the first moment of dN_{it}^D will be important for determining the *level* of asset prices. In normal times, dividends grow at an expected rate of g_{iD} , but if there is a disaster, the stock dividend experiences an additional shock, F_{it} . $F_{it} > 0$ captures how much dividends recover during a disaster. For instance, if $F_{it} = 0.9$, then dividends drop by 10% when a disaster arrives. Firms with high values of F_{it} therefore perform better during a crisis. The important thing to note at this juncture is that, inside of the model, firm dividend growth and aggregate consumption growth are both exposed to the same aggregate “jump” process, J_t . This is the sense in which disasters are macroeconomic — they impact all firms in the economy at the same time.

In order to compute equilibrium stock prices, we must put more structure on the arrival rate of disasters (p_t), the recovery rate of disasters at the macroeconomic level (B_t), and the dividend recovery rate during a disaster (F_{it}). One way to accomplish this is to model the dynamics of p_t , B_t , and F_{it} separately. Wachter (2013) is one example of this approach, as she models p_t as a Cox, Ingersoll, and Ross (1985) process. Alternatively, Gabaix (2012) models the dynamics of p_t , B_t , and F_{it} jointly, which I adopt in continuous time here. I do so through a variable called resilience, H_{it} , which summarizes how a firm is expected to perform if a disaster occurs over the

next instant, dt . I define stock i 's resilience exactly as in Equation (31) in Appendix C; however, in Appendix C I also show that a very good approximation that is more useful for understanding the underlying economics in the model is:

$$H_{it} = p_t \mathbb{E}_{t^-} [B_t^{-\gamma} (F_{it} + (\chi - 1)B_t) - \chi] + \frac{\sigma_C^2}{2} \gamma(1 + \gamma)$$

It is easy to see that firms that are expected to perform well in a disaster (i.e. firms whose dividends recover a lot) have high values of F_{it} , and thus are high resilience stocks. For ease of notation, I drop the limit notation of \mathbb{E}_{t^-} and replace it by \mathbb{E}_t , except when it is necessary to distinguish the two.

Resilience itself is modeled as a linearity generating (LG) process as in Gabaix (2009). Linearity generating processes behave, in an approximate sense, like an AR(1) process. The distinction between an LG process and an AR(1) process is subtle; however, LG processes are useful in economics because they deliver linear expressions for equilibrium price-dividend ratios without needing log-linearizations.⁷ As is common in macroeconomics, I further decompose resilience into a long-run mean, H_{i*} , and deviations from that mean, \hat{H}_{it} :

$$H_{it} = H_{i*} + \hat{H}_{it}$$

In turn, firm i 's resilience (deviations from mean) evolves according to the following stochastic

⁷Wachter (2013) and Seo and Wachter (2015) are examples of rare disasters models that uses log-linearizations. Indeed, in order to ensure model tractability, these authors adopt an IES of 1.

differential equation:

$$\begin{aligned} d\hat{H}_{it} &= -\phi_{H,i}\hat{H}_{it}dt - \hat{H}_{it}^2dt + dN_{it}^H \\ &\approx -\phi_{H,i}\hat{H}_{it}dt + dN_{it}^H \end{aligned}$$

where N_{it}^H is a martingale. It is easy to now see the relationship between an AR(1) process and an LG process. The term $-\phi_{H,i}\hat{H}_{it}dt$ would be present in the continuous time analogue of an AR(1) process, and so $\phi_{H,i}$ determines the speed of mean reversion for the resiliency process. It is sensible that resilience should behave similarly to an AR(1) process — if a firm has high resilience today, it is likely that resilience tomorrow is also high. The “extra” term, $-\hat{H}_{it}^2dt$, is unique to the LG class, and in practice is quite small because \hat{H}_{it} captures deviations from the mean of resilience.⁸ With the primitives of the model specified, I now turn to analyzing the equilibrium of this economy.

2.2.2 Equilibrium

In Appendices A and C, I derive the following equilibrium quantities: (i) the riskfree rate r_{ft} ; (ii) the wealth-consumption ratio; (iv) the price-dividend ratio of a claim to an arbitrary stock; (v) equilibrium returns under the physical measure, and (vi) equilibrium returns under the risk-neutral measure. For much of my empirical work, I will focus on the equilibrium returns of an arbitrary stock i under the risk-neutral measure. In the model, risk-neutral returns take a

⁸As noted in Gabaix (2009), there are restrictions on the variance of the noise process dN_{it}^H that ensure prices remain positive. In Appendix C.3, I provide one example of a process that satisfies the restrictions and is also internally consistent with the rest of the model.

simple form:

$$\frac{dP_{it}^*}{P_{it}^*} = r_{ft}dt + \sigma_{it}dW_{it}^* - p_t\mathbb{E}_t[B_t^{-\gamma}(F_{it}-1)]dt + (F_{it}-1)dJ_t^*$$

where W_{it}^* is a Brownian motion under the risk-neutral measure and σ_{it} is defined in Equation (35). dJ_t^* is the disaster process under the risk-neutral measure and it arrives at a risk-neutral rate of $p_t^* := p_t\mathbb{E}_t[B_t^{-\gamma}]$. It is useful at this point to compare p_t^* to p_t . When disasters are bad for the economy ($\mathbb{E}_t[B_t] < 1$), then the risk-neutral likelihood of a crisis is higher than the actual likelihood, i.e. $p_t^* > p_t$. This occurs because the risk-averse agent prices assets as if disasters are more likely than they actually are in nature.

Armed with risk-neutral equilibrium returns, it is also straightforward to work out that the jump-risk measure $\mathbb{D}_{it}(\tau)$ from Section 2.1 maps to the following quantity *inside of the model*:

$$\mathbb{D}_{it}(\tau) = \Psi_i \times p_t^*(\tau) \tag{2}$$

where $p_t^*(\tau) := \mathbb{E}_t^* \left[\int_t^{t+\tau} p_s B_s^{-\gamma} ds \right]$ is the risk-neutral expected probability of a disaster from time t to $t + \tau$. Again, the exact expression for Ψ_i is given in Equation (47) of Appendix C.6 and it depends on the higher moments of F_{it} . In deriving the mapping between the model and $\mathbb{D}_{it}(\tau)$, the working assumption that I have made is that the distribution of disasters F_{it} is time-invariant and independent of the size of the aggregate recovery rate, B_t . These assumptions are made for the sake of parsimony, and I discuss how they impact the interpretation of my results in Section 3.2. Often times, I will drop the argument τ and simply write p_t^* ; in this case, the reader can assume that $\tau = 30$ days, which corresponds to the horizon that I study in my empirical analysis.

The important takeaway at this juncture is that at each time t , rare disasters models imply

that all cross-sectional variation in $\mathbb{D}_{it}(\tau)$ is driven by variation in Ψ_i . This holds true only because $p_i^*(\tau)$ is common to all firms in the economy. Indeed, we saw in Section 2.1 that when each firm has its idiosyncratic probability of a jump $p_{it}^*(\tau)$, cross-sectional variation in $\mathbb{D}_{it}(\tau)$ was instead driven by variation in both Ψ_i and $p_{it}^*(\tau)$. In other words, the rare disasters model restricts the risk-neutral probability of a jump to be the same across all firms, $p_{it}^*(\tau) = p_i^*(\tau)$.

Econometrically, this means that under the rare disasters hypothesis, a panel of $\mathbb{D}_{it}(\tau)$ should obey a latent factor structure that is linear in a single factor. To see this more forcefully, suppose that my estimate of $\mathbb{D}_{it}(\tau)$ is denoted by $\widehat{\mathbb{D}}_{it}(\tau)$. Further suppose that the true $\mathbb{D}_{it}(\tau)$ is measured with some idiosyncratic measurement error, e_{it} , such that:

$$\begin{aligned}\widehat{\mathbb{D}}_{it}(\tau) &= \mathbb{D}_{it}(\tau) + e_{it} \\ &= \Psi_i \times p_i^*(\tau) + e_{it}\end{aligned}\tag{3}$$

The idiosyncratic measurement error can be interpreted as pure noise in the option prices or as unmodeled idiosyncratic jump risk.⁹ Because we do not observe Ψ_i or $p_i^*(\tau)$, Equation (3) is a latent factor model. As is well-known in latent factor analysis, the loadings (Ψ_i) and the factor ($p_i^*(\tau)$) are not separably identifiable. Thus, we can estimate time-series variation in $p_i^*(\tau)$. Nonetheless, this procedure cannot pin down the mean of $p_i^*(\tau)$, though it would certainly be useful to do so given that the academic debate about these models has centered around how often (and severe) disasters actually happen. Still, the critical assumption underlying the entire class of rare disaster models is that $p_i^*(\tau)$ is common to all firms, and this easily testable hypothesis forms the basis for the empirical work in this paper.¹⁰ Finally, to avoid notional confusion,

⁹For instance, consider a case where each firm is exposed to a systematic jump risk factor and an idiosyncratic jump risk factor. In this case, it is possible to show that $\mathbb{D}_{it}(\tau)$ will be linear in both jump-risk factors. In turn, the idiosyncratic jump risk factor will just be absorbed into the noise term e_{it} .

¹⁰For the entire class of rare disaster models, the pricing kernel is subject to large and rare jumps, corresponding

I will henceforth not distinguish between $\mathbb{D}_{it}(\tau)$ and $\widehat{\mathbb{D}}_{it}(\tau)$; however, implicit in my remaining analysis is the latent factor structure and measurement errors that are defined in Equation (3).

3 An Estimate of the Risk-Neutral Probability of Disaster

In this section, I build measures of $\mathbb{D}_{it}(\tau)$ using a large cross section of option prices, thereby allowing me to test the model's central prediction that $\mathbb{D}_{it}(\tau)$ should be governed by a single factor. This procedure also provides me a way to extract an estimate of the risk-neutral probability of a disaster, p_t^* . Finally, in the latter part of the section I explore the predictions of general equilibrium models of disasters by relating p_t^* to future economic activity.

3.1 Construction and Data Description

As alluded to earlier, each jump risk measure $\mathbb{D}_{it}(\tau)$ is really just a portfolio of call and put options written on firm i , as of date t . Thus, to construct each $\mathbb{D}_{it}(\tau)$ measure, I proceed in two steps: (i) collect option prices (across strike prices) of maturity τ on date t ; (ii) combine call and puts of different strikes according to the portfolio weights in [Du and Kapadia \(2013\)](#). $\mathbb{D}_{it}(\tau)$ is simply the resulting price of the portfolio. The exact weights take a fairly complex form, and explicit formulas can be found in [Appendix B.1](#).

The data used to construct the panel of $\mathbb{D}_{it}(\tau)$ comes from OptionMetrics. For each firm and each date, I collect only out-of-the-money (OTM) options that meet the following criteria: (i) open interest is strictly positive, (ii) best offer minus best bid is strictly positive, (iii) best bid is strictly positive, (iv) no special settlement; (iv) the maturity of the option is greater than 7 days;

to large increases in marginal utility. If this is the case, then risk-neutral returns for all firms in the economy will be exposed to a common jump process. Thus, the arguments in this section are not necessarily tied to this specific disaster risk model.

and (vi) implied volatility is not missing. This screening procedure requires me to merge raw option prices with underlying security information, which I obtain from OptionMetrics as well.

The construction of my final measure proceeds in a two steps. I first compute a $\mathbb{D}_{it}(\tau)$ according to Equation (20) for each available option maturity on date t . Next, I compute a $\mathbb{D}_{it}(\tau = 30 \text{ days})$ by linearly interpolating between the $\mathbb{D}_{it}(\tau)$ for maturities that straddle 30 days. For the remainder of the paper and with some abuse of notation, I use \mathbb{D}_{it} to denote the 30-day implied measure. A complete description of this procedure can be found in Appendix B.

The end result is a panel of $\mathbb{D}_{it}(\tau = 30 \text{ days})$ for 6,762,860 firm-day pairs. For each firm, I also construct a monthly measure by averaging within each month. For example, if the current month is January 1997, then the monthly measure of \mathbb{D}_{it} effectively represents the risk-neutral expected jump risk of firm i during February 1997. When aggregating to a monthly measure, I require that each firm has at least 70% of its daily observations available for that month.

3.2 Results

I begin my empirical analysis by conducting principal component analysis on the panel of \mathbb{D}_{it} . I analyze the monthly \mathbb{D}_{it} because it allows me to study a larger cross section of firms, and it also smooths daily measurement noise. To be included in the analysis, a firm must have its entire time-series available for the period January 1996 to May 2015, leaving 242 firms in total.¹¹

Figure 1 indicates that there is a strong one-factor structure driving time-series variation in \mathbb{D}_{it} across firms. The first principal component is dominant, explaining nearly 50% of the variation across \mathbb{D}_{it} . The remaining principal components are much weaker. The second principal

¹¹Specifically, a firm must have at least 70% of its daily observation in a month to be included. Principal component analysis is conducted on the correlation matrix of monthly \mathbb{D}_{it} . This analysis applies to the set of firms with at least 80% of its monthly observations. To generate the time-series of the factor, I fill missing data in with mean values. The results are not sensitive to these choices.

component captures a little over 10% of variation and the third component accounts for less than 5%. The factor structure inherent in \mathbb{D}_{it} is consistent with the model's core assumption of a single disaster event driving asset prices for all firms in the economy. I therefore interpret the first extracted principal component as p_t^* . Appendix B.6 contains a more detailed mapping of the extracted factor to model quantities.

Figure 2 displays the time-series of the risk-neutral probability of disaster. It is clear that p_t^* is very persistent (monthly AR(1) coefficient of 0.91), which is what we would expect if it reflects the likelihood of macroeconomic crises. It also seems reasonable to view p_t^* as the probability of disaster based on its evolution through time, as its peaks correspond to some well-known crisis episodes. For instance, p_t^* spikes after the terrorist attack on the United States in September 2001. Similarly, it rises in July 2002 in response to WorldCom filing for bankruptcy (at the time this was the largest corporate insolvency ever), increased turmoil in the Gaza Strip, and an earthquake in Germany.

The focal point of rare disasters models is large drops in aggregate consumption, and wars are often cited as an example of these deep recessions. Consistent with this notion, a prolonged period of high disaster probability occurs at the onset of the U.S.-Iraq War. Moreover, the technology bubble that burst at the end of the 1990s does not seem to have the same effect on disaster likelihood as the war, indicating that recessions don't drive all movements in this time series.

The period from 2002 to 2007 corresponds to a relatively low value for p_t^* , with disaster risk rising at the beginning of the U.S. financial crisis in August 2007. The massive increase in p_t^* in October 2008 followed the bankruptcy of Lehman Brothers in September 2008 and a series of global stock market crashes at the end of October 2008. The maximum for the sample occurred in November 2008, the same month that the U.S. government injected \$40 billion into AIG, a

proposal to bailout U.S. automakers failed in the U.S. Senate, and the government injected \$20 billion into Citigroup. Overall, the six largest months for p_t^* occur during the pinnacle of the U.S. financial crisis from October 2008 to March 2009.

There is another spike in May 2010 following the so-called “flash crash” and the spread of the European debt crisis as Spain, Portugal, and Greece all saw their economic conditions worsen significantly. The last major rise in disaster probability occurs in August 2011, when there were concerns about the European sovereign debt crisis worsening, as well as worries over a potential downgrade of the creditworthiness of U.S. debt.

In Table 1, I compute the correlation of p_t^* with other popular financial market indicators, both in levels and first differences. Unsurprisingly, periods of high p_t^* coincide with low prices, as evidenced by the -23% correlation in levels (-50% in first differences) with the price-dividend ratio of the S&P 500. In addition, p_t^* appears to be modestly correlated with the TED spread (50% in levels, 35% in first differences), indicating that a rise in the risk-neutral probability of a disaster occurs when financial market liquidity dries up.¹² Finally, changes in p_t^* display a modest negative correlation with innovations to financial intermediary leverage (Adrian et al. (2014)), suggesting that financial firms deleverage as disaster risk increases. In sum, the evolution of p_t^* and its correlation with other financial market variables are consistent with the interpretation that p_t^* measures variation in the risk-neutral probability of a disaster.

Alternative Interpretation I have interpreted the common factor driving movements in \mathbb{D}_{it} as measuring changes in the risk-neutral probability of a disaster, p_t^* . An equally plausible interpretation is that I’m picking up common movements in the severity of disasters, and I certainly

¹²Unsurprisingly, p_t^* is highly correlated with disaster risk \mathbb{D}_{SPX} of the S&P 500 index (88% in levels, 69% in first differences). However, a closer inspection of the two series reveals the advantage of estimating p_t^* using principal component analysis instead of directly using \mathbb{D}_{SPX} ; the mean-to-variance ratio of p_t^* is about 0.79 whereas the same ratio for \mathbb{D}_{SPX} is 0.53. Insofar as this metric captures the signal-to-noise ratio of a series, then p_t^* is measured with 50% more precision than \mathbb{D}_{SPX} .

can't rule this out, nor do I really need to. To see why, consider the most general case where aggregate disaster probability (p_t) and recovery (B_t) are time-varying, and so is the distribution of firm-level disaster recovery (F_{it}). Without any loss of generality, we can assume that Ψ_{it} — a function of the distribution of F_{it} — can be decomposed into a common component, Ψ_t , and an idiosyncratic component, ξ_{it} . In this case it is easy to show that the jump-risk measure \mathbb{D}_{it} corresponds to:

$$\mathbb{D}_{it} = p_t^* \Psi_t + p_t^* \xi_{it}$$

which implies that there should be one very strong factor driving \mathbb{D}_{it} , and potentially a weaker second factor, depending on the variability of ξ_{it} . Because the data broadly supports this general case, my focus is on the first extracted factor. The larger point here is that the factor structure of the \mathbb{D}_{it} tells us something about common time-series variation in the probability and/or severity of a jump across all firms. Rare disaster models make strong predictions for what this common variation should look like, and this is really what I'm taking advantage of in this paper. As we will also see, the distinction between a common time-varying probability versus severity of a disaster is not particularly important for my remaining empirical findings.

3.3 p_t^* and Future Economic Activity

If movements in p_t^* are in fact capturing macroeconomic disaster risk, then it is natural to expect p_t^* to relate to broader macroeconomic activity. Indeed, [Gourio \(2012\)](#) builds a general equilibrium model where disasters impact both total factor productivity (TFP) and the aggregate capital stock. A key insight of this model is that an increase in disaster likelihood — not necessarily a realization of a disaster — can itself have a profound impact on the macroeconomy. For instance, when the likelihood of a disaster increases, firms reduce investment because

physical capital becomes riskier. In addition, an imminent disaster reduces the return on savings, meaning unemployment increases because work today is less attractive. [Kilic and Wachter \(2015\)](#) also build a model where disaster risk increases unemployment by lowering job-creation incentives. Like in [Gourio \(2012\)](#), future unemployment responds to disaster risk itself without needing an actual shock to labor productivity.

To empirically test whether an increase in the probability of a disaster leads to contractions in economic activity, I conduct simple forecasting regressions of the following form:

$$Y_{t+1} = c + \sum_{i=1}^4 \gamma_i Y_{t+1-i} + \boldsymbol{\phi}' \mathbf{X}_t + \beta p_t^* + u_{t+1}$$

where Y_{t+1} is a measure of economic activity and \mathbf{X}_t is a vector of competing forecasting variables. The macroeconomic activity measures that I use are as follows: the (i) change in the unemployment rate; (ii) percentage change in the Industrial Production Index; (iii) change in total industry capacity utilization; and (iv) percentage change in the ISM Manufacturing Inventories Index.¹³ Following [Gilchrist and Zakrajsek \(2012\)](#), I use standard financial market variables as my competing forecasting variables. Specifically, \mathbf{X}_t includes the real federal funds rate, the term spread from U.S. treasury yields (10-year minus 3-month constant maturity yield), and the excess return of the CRSP-value weighted index.¹⁴

The forecasting regression results in column (1) of Panels A-D in [Table 2](#) indicate that p_t^* is statistically a highly significant predictor of future economic conditions. Also consistent with theory, an increase in p_t^* forecasts an economically meaningful decline in activity. For instance, a one standard deviation increase in p_t^* translates to an increase in unemployment of 0.065

¹³The unemployment rate comes from the Bureau of Labor Statistics (series id = LNS1400000). The industrial production index (series id = INDPRO), total capacity utilization (series id = TCU), ISM Manufacturing: Inventories Index (series id = NAPMII) were obtained from the St. Louis Federal Reserve FRED database.

¹⁴The real federal funds rate is computed as the average monthly effective federal funds rate within month t , minus the growth rate of the PCE price index between months $t - 1$ and $t - 13$.

percentage points (e.g. an increase from 5.8% unemployment to 5.865%). The U.S. labor force is roughly 160 million people, so this translates to about 100,000 lost jobs per month. Moreover, panels B-D in Table 2 suggest that investment does indeed fall as disasters become more likely. A rise in p_t^* leads to fairly large drops in total capacity utilization, industrial production, and the manufacturing inventories index.

The results in column (1) of Panels A-D in Table 2 also indicate that, in the short-run, the macroeconomy almost always responds more movements in p_t^* relative to other popular financial market indicators. For instance, a one standard deviation increase in p_t^* lowers industrial production by 0.16% per month. This may not appear to be a particularly large response, yet industrial production responds even less to the term spread, the real federal funds rate, and the return on the aggregate stock market. Overall then, the forecasting exercise demonstrates that p_t^* contains important information about future economic downturns — information that is also not contained in standard financial market indicators.

Additional Robustness It is natural to wonder whether these forecasting results are driven a single data point, namely the 2008 financial crisis. To explore this idea further, I rerun the forecasting regressions, but exclude the NBER recession years from the crisis (December 2007 to June 2009). Column (3) of Panels A-D in Table 2 indicates that the forecasting results generally go through when excluding the crisis. The results are somewhat weaker though, as evidenced by the lower forecasting R^2 . Nonetheless, when forecasting on future unemployment, industrial production, and total capacity utilization, the magnitude and significance of the point estimates on p_t^* are also fairly robust to excluding the crisis.

To further highlight the unique informational content of p_t^* , in Column (2) of Panels A-D in Table 2, I also include the VIX index in the forecasting regressions. Column (4) repeats the

forecasting regressions with the VIX, but excludes the crisis period. The VIX is a natural benchmark in this setting because it also derives from option prices. Indeed, p_t^* is 85% correlated with the VIX in levels and 63% in first differences. However, this is not really that surprising; as shown by [Du and Kapadia \(2013\)](#) and [Martin \(2013\)](#), the VIX is a biased estimate of risk-neutral volatility, and the bias comes precisely from the jump component of risk-neutral returns.¹⁵ Put differently, the VIX contains information on both disasters and normal-times volatility, whereas p_t^* measures only time-varying disaster risk.

In a majority of cases, the inclusion of the VIX does not drive out the forecasting ability of p_t^* . In all specifications, both with and without the crisis, the point estimate on p_t^* is basically unchanged when including the VIX. Moreover, in five out of eight cases, p_t^* remains statistically significant at a 10% confidence level. The latter finding is somewhat surprising, given that one would expect the standard errors in the regression to rise because the VIX and p_t^* are highly correlated. Moreover, the VIX has economically small point estimates in nearly all specifications. In total, it therefore seems reasonable — and consistent with rare disaster models — to conclude that p_t^* is a unique and robust forecaster of future economic activity.

4 Disaster Risk and The Cross Section of Equity Returns

Thus far, much of the empirical analysis has centered around the model's predictions for the factor structure in option-implied jump measures and the estimation of p_t^* , a single measure of aggregate disaster risk. Still, the model makes strong predictions for the *cross section* of equity returns as well. The easiest way to see this is from Theorem 2 in Appendix A, which states that the expected excess return $\mu_{it} - r_{ft}$ for an arbitrary stock i is decreasing in the expected amount

¹⁵In fact, the \mathbb{D}_{it} measure constructed earlier takes advantage of this bias to isolate jump risk from option prices.

that it recovers during a crisis, F_i :

$$\mu_{it} - r_{ft} = p_t \mathbb{E}_t [B_t^{-\gamma} (1 - F_{it})]$$

This result makes intuitive sense because if a firm is expected to recover a lot during a crisis (high F_{it}), it should command a low risk premium. To this end, if given an estimate of $\mathbb{E}_t [F_{it}]$ at each point in time, one could sort stocks based on $\mathbb{E}_t [F_{it}]$.¹⁶ In turn, the model predicts that stocks with high values of $\mathbb{E}_t [F_{it}]$ should earn low average returns, and vice versa. The following proposition provides a simple way to operationalize this using basic OLS regressions and my estimated p_t^* .

Proposition 1. *Within the model, consider an OLS regression of realized returns for firm i , r_t^i , on changes in the risk-neutral probability of disaster:*

$$r_t^i = a + \beta_i \Delta p_t^* + e_t^i$$

Assume B_t and F_{it} are independent and that expectations of these variables are constant over small increments. Then β_i is increasing in $\mathbb{E}_t [F_{it}]$, i.e. $\partial \beta_i / \partial \mathbb{E}_t [F_{it}] > 0$. This is true even if the true mean of p_t^ is unknown. Thus, firms with higher β_i should earn lower average returns.*

Proof. See Appendix C.7. □

The intuition of Proposition 1 is straightforward: stocks that move in tandem with p_t^* (positive β_i firms) are valuable hedges against disasters. These stocks do not command high risk premiums. Furthermore, even when the mean of p_t^* is not known — the case for the p_t^* estimated in the previous section — the *relative* rankings based on β_i do not change. This is because the

¹⁶Clearly, this assumes that B_t and F_{it} are independent.

dispersion in β_i at each point in time is unaffected by the mean of p_i^* .¹⁷ Henceforth, I call β_i a firm’s “disaster- β ” in order to distinguish it from the standard CAPM β .

It is also important to recognize that disaster- β ’s are a good way to detect *priced* disaster risk in more general settings. For instance, one could easily add idiosyncratic jumps to the firm dividend process in the model. It is straightforward to work out that the model-implied \mathbb{D}_i will be additive in the aggregate risk-neutral jump process, p_i^* , and the idiosyncratic jump process for each firm. In this case, \mathbb{D}_i will continue to obey a single-factor structure, with the single factor still representing p_i^* .¹⁸ Thus, disaster- β ’s will still pick up on the priced disaster component of each firm’s equity.

4.0.1 Empirical Implementation

In lieu of Proposition 1, I apply the following procedure. For each firm, I compute a weekly \mathbb{D}_i by taking within-week averages of that firm’s daily \mathbb{D}_i . At the end of each month m , I extract the first principal component of the previous year’s estimates of weekly \mathbb{D}_i . I call this my estimate of p_i^* for month m . Next, I regress firm i ’s weekly equity returns on Δp_i^* using data over the previous year. This delivers an estimate of firm i ’s disaster- β . Finally, at the end of month m , I sort equities into quintiles based on their estimated β_i . I hold these equal-weighted portfolios for one month, then repeat the entire procedure. The sample of stocks that I examine comes from CRSP. I consider only common stocks (CRSP SHRCD 10 or 11) that are listed on the NYSE, NYSE Amex, or NASDAQ. To avoid penny stocks, I also exclude stocks with a share price of less than \$1 at the time of portfolio formation.

There are two aspects of this approach that are worth noting. First, the trading strategy is fully implementable in real-time: p_i^* and the β_i ’s are all estimated using only data that was

¹⁷i.e. Scaling p_i^* by a positive constant c simply scales all β_i by c . It does not change relative rankings.

¹⁸i.e. the idiosyncratic jump will be absorbed into the noise term in the Equation (3).

available at the end of month m . To further ensure tradability, I measure disaster- β 's on the last business day of each month, but I form the portfolios at the end of the next business day.¹⁹ Second, I am able to discern which stocks are disaster risky, even if these stocks do not have traded options. This is very useful because it allows me to examine the pricing implications of the disaster model for a much broader set of equities.

To give an example of my empirical strategy, suppose I am standing at the end of January 1997. I first construct a weekly p_t^* series using option data from January 1996-January 1997. Next, for each firm in my CRSP sample, I regress weekly returns from January 1996 to January 1997 on Δp_t^* to get β_i . I then put firms with the largest β_i 's into one portfolio, and so forth. I form the portfolios on the first day in February 1997, hold them until the end of the month, and then repeat the procedure.

Finally, following [Frazzini and Pedersen \(2014\)](#), I use a simple approach to estimating each β_i :

$$\hat{\beta}_i = \frac{\rho(r_t^i, \Delta p_t^*) \times \sigma(r_t^i)}{\sigma(\Delta p_t^*)}$$

where $\sigma(x)$ is the estimated volatility of series x , and $\rho(y, x)$ is the estimated correlation between x and y . A more detailed description of how I create these equity portfolios can also be found in [Appendix B](#).

4.1 Results

Table 3 provides summary statistics for the portfolios formed based on disaster risk. As expected, the raw average excess return is monotonically increasing in disaster risk. High disaster risk stocks (portfolio 5) earn 12.26% in annualized excess returns whereas low disaster risk firms

¹⁹This means the market values used for value-weighting are as of the beginning of the following month. Similarly, realized returns are computed from the first of the month to the end of the month.

earn only 5.08%. Row 4 of Table 3 also suggests that disaster risky stocks tend to be smaller, with an average market capitalization of about \$1.47 billion.

The defining feature of disaster risky stocks is their sensitivity to movements in p_t^* . As is evident from the fifth row of Table 3, disaster risky stocks display a negative contemporaneous relationship with changes in p_t^* , whereas low disaster risk stocks have positive comovement with p_t^* . The dispersion in estimate disaster- β 's across quintiles is also quite strong, ranging from 14 to -12. Later in Section 4.2, I use a calibrated version of the model to interpret these estimated β 's further.

Table 4 displays the risk-adjusted returns for each portfolio, which are computed using exposure to the three Fama-French factors and a standard momentum portfolio.²⁰ Four-factor alphas increase monotonically when moving from low disaster risk to high disaster risk. A portfolio that is long disaster risky stocks and short low disaster risk stocks earns 7.56% in annualized alpha (t -statistic of 2.35). Though the point estimate on the alpha for low disaster- β stocks is measured imprecisely, the negative sign indicates that investors are actually willing to pay a premium for firms that hedge their exposure to macroeconomic disasters.

Interestingly, the CAPM- β is relatively flat when moving from low to high disaster risk (row 3 of Table 4), suggesting that disaster- β contains information that is not captured by simple exposure to the aggregate stock market. This finding is consistent with Gabaix (2012)'s argument that CAPM- β may poorly measure disaster exposure due to a lack of realizations in the data. This observation also highlights the usefulness of deriving disaster- β based on exposure to an option-based measure, p_t^* : one does not necessarily require a disaster to be realized in order to measure disaster risk.

Rows 5 through 10 summarize the exposure of the disaster-sorted portfolios to the Fama-

²⁰All risk factors were obtained from Ken French's website.

French size and value factors, as well as the standard momentum factor. Because there is not a great deal of dispersion in the loadings on the Fama-French size factor, the long-short portfolio (5-1) has only a slightly negative exposure to size. In a similar vein, the exposure of the 5-1 portfolio to momentum is nearly zero. On the other hand, it appears that disaster risky stocks have some relation to value. The loadings on the value factor generally increase when moving from low to high disaster risk. Nonetheless, the magnitudes of the loadings are not particularly high, as the 5-1 portfolio has a value- β of only about 0.5. Thus, exposure to standard risk factors does not seem to explain the dispersion of excess returns generated by differential disaster risk.

The last row of Table 4 presents the test that the alphas of portfolios 1-5 are jointly zero. As indicated by a p -value of 1.3%, the null hypothesis of this test is strongly rejected. In summary, these results suggest that disaster risk is priced in the cross section of U.S. equities. Consistent with the model, disaster risky stocks require high excess returns and the premium for bearing disaster risk appears to be around 7.6% per year.

Robustness

Tables 5 - 7 present various robustness tests of the claim that disaster risk is priced in the cross section of equities. Table 5 constructs value-weighted, as opposed to equal-weighted, portfolios based on disaster- β . The patterns are generally the same as the equal-weighted portfolios, with the long-short portfolio earning an even higher alpha of 11.8% (t -statistic of 3.11). Like with the equal-weighted case, the 5-1 portfolio does not have large loadings on the market, size, or momentum portfolios. There still appears to be a relation between disaster risk and value investing. However, this relationship is not enough to explain the dispersion in alphas generated by differential loading on disaster risk. The joint test that all portfolio alphas are equal to zero is also strongly rejected.

In Tables 6 and 7 display the results of the portfolio sorts (both equal and value-weighted) when disaster- β is computed using daily data, as opposed to weekly data.²¹ The results are only slightly weaker, but there is a natural explanation for this. The estimate for p_t^* is much noisier at a daily level and stocks with thin trading have no daily price changes. The combination of these two issues adds measurement error to the estimation of disaster- β 's, thereby creating much less dispersion because disaster- β 's shrink towards zero. Nonetheless, from the first row of Tables 6 and 7, it is still the case that high disaster risk stocks earn higher risk-adjusted returns than low disaster risk stocks. Moreover, when forming portfolios on a value-weighted basis, the 5-1 earns a four-factor alpha of about 10% and this alpha is significantly different from zero in a statistical sense. The joint test that all of the alphas are zero is also rejected at a 1% confidence level when value-weighting.

4.2 Two Calibration Exercises

4.2.1 How Big Should Disaster- β 's Be?

For a calibrated version of the model, how sensitive should returns be to changes in p_t^* ? To simplify the calibration, I assume that B_t and F_{it} are constants. Further, I assume that p_t is constant within each month. Together, these two assumptions deliver a simple expression for the model-implied disaster- β_i .

Corollary 1. *Suppose B_t and F_{it} are constant, and p_t is constant within each month. Consider the regression:*

$$r_{it} = a_i + \beta_i \Delta p_t^* + e_{it}$$

²¹To compute a daily disaster- β I estimate: $r_{it} = a_i + \sum_{k=0}^2 \beta_{ik} \Delta p_{t-k}^* + e_{it}$, at a daily frequency. As is common in the literature, the lags account for asynchronous trading. The disaster- β for firm i is then $\sum_{k=0}^2 \beta_{ik}$. The regression uses the previous quarter's worth of data.

where p_i^* is the estimated risk-neutral disaster likelihood in Equation (21). Then β_i is a function of model primitives:

$$\beta_i = \frac{B^{-\gamma} (F_i + (\chi - 1)B) - \chi}{B^{-\gamma} (\delta_i + \phi_{H,i}) \times \bar{\Psi}}$$

where $\bar{\Psi}$ corresponds to the average $\Psi_i = 2 \times \{1 + \ln(F_i) + [\ln(F_i)]^2 / 2 - F_i\}$ in the economy.

Proof. See Appendix C.8 □

Based on Corollary 1, I conduct a calibration exercise. Most of the parameters are close to Barro (2006) and Gabaix (2012). For instance, I set $B = 0.60$, corresponding to a 40% drop in consumption during a disaster. Similarly, I compute $\bar{\Psi}$ by using the average \bar{F} across firms. I set $\bar{F} = B = 0.6$, indicating that the aggregate stock market falls 40% in a crisis. This number is roughly consistent with the recent experience during the 2008 U.S. financial crisis.²² I use a modest risk-aversion of $\gamma = 4$, which in the previously mentioned studies is large enough to match the unconditional average equity risk premium. Finally, I assume the intertemporal elasticity of consumption (IES) is equal to 2, which produces $\chi = -1/6$.

Figure 3 generates a range of β_i for different values of firm recovery, F_i . All of the parameters used in the calibration are listed along the top of the figure. As detailed in Proposition 1, β_i is monotonically increasing in F_i . Loosely speaking, this is why sorting on β_i is equivalent to sorting on F_i . Additionally, for reasonable values of $F_i \in [0.5, 0.8]$, β_i ranges from about -25 to 15. The range of these model-implied disaster- β 's roughly matches their data counterparts, as reported in row 5 of Table 3.

To hone the calibration further, in Table 8 I empirically estimate the disaster- β of the S&P 500 and find $\beta_{SP} \approx -6.4$. Under the current calibration exercise, the model implies that $\beta_{SP} =$

²²Alternatively, one could easily incorporate uncertainty about the severity of disaster. For instance, Barro and Jin (2011) estimate a power law distribution to the size of disasters. When the size of disasters is fat tailed, then the average disaster size can be much smaller than 40% and still generate large risk premia. As this calibration is largely illustrative, I use the simpler assumption of a constant disaster severity.

–11.²³ By no means is this comparison a formal test, but the model does line up reasonably well with the data along this particular dimension. More broadly, the empirical fact that the market falls in value as p_t^* rises implies an IES greater than one. This point is also argued by, among others, Nakamura, Steinsson, Barro, and Ursua (2013).

Assuming a constant firm recovery rate F_{it} was useful in the current calibration exercise to get a sense of average disaster- β 's. Still, how likely is it that F_{it} is constant? The results in Table 8 also shed some light on this question. From Table 8, variation in p_t^* can explain only about 25% of the variation in aggregate stock returns. Gabaix (2012) finds a similar, albeit slightly smaller value of 13%.²⁴ In the context of the model, the fact that time-series variation in p_t^* accounts for a relatively small amount of return variation implies that the remaining movements are driven by F_{it} . It appears then that F_{it} must vary through time, which is a one reason I estimate disaster- β 's using rolling regressions when performing portfolio sorts.

4.2.2 Do Disaster Risky Stocks Fall Enough to Justify Their High Returns?

The results in this section suggest that disaster risk is a priced risk for the cross section of equity returns. High disaster risk stocks earn high returns precisely because they are the most liable to lose value during macroeconomic crises. To assess this statement within the model, I ask a simple question: how much does the model imply disaster risky stocks must lose in order to justify their returns? Then, I compare the model-implied loss rates to what we actually see during the 2008 U.S. financial crisis, arguably classified as a macroeconomic disaster.

As a loose evaluation, I focus on the unconditional average (risk-adjusted) excess returns of the disaster risk portfolios from earlier in the section. I infer a model-implied F_i for each portfo-

²³Implicitly this means I set $F_{SP} = \bar{F} = 0.6$.

²⁴The 13% comes via a regression of changes in the aggregate price-dividend ratio on changes in the real riskless rate. The idea in that setting is that movements in the riskless rate reflect only p_t .

lio based on the equilibrium relationship $\mu_{it} - r_{ft} = pB^{-\gamma}(1 - F_i)$. As in the previous calibration exercise, I set the recovery of consumption during a disaster to $B = 0.60$ and risk-aversion $\gamma = 4$. Following [Nakamura, Steinsson, Barro, and Ursua \(2013\)](#), I also set the probability of a disaster $p = 0.028$. For the sake of comparison, I compute the realized recovery of each portfolio as (one minus) the minimum 12-month return over the period December 2007 to June 2009. This period corresponds to the NBER recession that encompassed the financial crisis.

Table 9 collects the results of this exercise. The data is loosely in the line with the model, though there are some interesting discrepancies. For portfolios 3-5, their realized recoveries during the crisis are higher than the model-implied recovery rates. For instance, the portfolio with the most disaster risk (portfolio 5) recovers about 49% of its value during the crisis, yet its 12.26% excess return implies that it should only recover 43% of its value. Nevertheless, there is a natural reason why realized recovery rates may be higher than model-implied recovery rates, at least for the more disaster risky stocks. During the financial crisis, the U.S. government conducted unprecedented interventions into capital markets. In turn, government bailouts likely put a lower bound on losses for disaster risky firms (e.g. financials), thereby inflating their realized recovery rates. [Nakamura, Steinsson, Barro, and Ursua \(2013\)](#) also argue that realized prices during disasters are likely to be distorted given that governments typically implement various forms of price controls during crisis episodes.

The more puzzling discrepancy between the model and the data is for the least disaster risky portfolio (portfolio 1). This portfolio recovered only 46% of its value whereas the model suggests it should have recovered 76%. Still, it is likely that this difference is driven by the fact that there is essentially one risk-factor in the model, perhaps suggesting that additional risk-factors would better match the data. Overall though, the magnitudes of the realized recovery rates are at least roughly consistent with those implied by the model.

5 Conclusion

This paper provides evidence of a single factor that drives the cross section of option-implied jump risk. I use a standard model of rare disasters to interpret this factor as the risk-neutral probability of a disaster, p_t^* . From 1996-2014, movements in p_t^* correspond to well-known crises, and also forecast declines in economic activity. In addition, I use p_t^* to show that disaster risk is priced in the cross section of U.S. equities: investors require a premium of 7.6% to hold disaster risky stocks.


A benefit of my approach is that I estimate time-series variation in p_t^* by starting from a non-parametric measure of option-implied tail risk. This is useful because it makes almost no assumptions on the data a priori, but the weakness is that I am not able to estimate the average likelihood of a crisis. A natural next step is to apply a structural model of disasters to the entire cross section of options. For instance, the model in [Seo and Wachter \(2015\)](#) applies to the aggregate stock market, but one could easily extend it to an arbitrary cross section of firms. Because the arrival rate of disasters is the same across firms (as shown in this paper), the number of additional parameters that needs to be estimated is small, effectively one for each firm. In addition, the cross section of options would add tremendous power in identifying the parameters of the disaster arrival process itself. This approach should also become more computationally tractable as computing power continues to increase.²⁵

On the theoretical side, an important outstanding question is what drives time-variation in the likelihood of a disaster in the first place? In this particular paper, I estimate the risk-neutral likelihood of a disaster, meaning its time-series movement can be driven by changes in the physical expectation of a disaster, changes in preferences, or both. In turn, what drives

²⁵[Santa-Clara and Yan \(2010\)](#) estimate a similar model using only S&P 500 options and find the average estimation time to be 2-4 hours. Using options across all strikes and all firms should make the problem more complex.


changes to, say, the physical likelihood of a disaster in general equilibrium? A potential answer could be found by linking the rare disasters literature to research on how financial intermediary activity impacts capital markets and the real economy. Furthermore, I have also worked out of a standard rational expectations framework, though it would be interesting to explore whether changes in p_t^* are driven by learning, sentiment, or other behavioral explanations.

Figure 1: Principal Component Analysis of $\widehat{\mathbb{D}}_i$ Across Firms




Notes: This figure displays the results of principal component analysis on the correlation matrix of firm disaster risk measures, $\widehat{\mathbb{D}}_i$, extracted from option prices. A monthly time-series of $\widehat{\mathbb{D}}_i$ is constructed by taking the within-month average of daily $\widehat{\mathbb{D}}_i$ for each firm i . A firm must have at least 70% of its daily observation in a month to be included. Principal component analysis is conducted on the correlation matrix of monthly $\widehat{\mathbb{D}}_i$. This analysis applies to the set of firms with at least 80% of its monthly observations. The date range is from January 1996 to May 2015.

Figure 2: Time-Series of p_t^*



Notes: This figure displays the first principal component from a panel of firm disaster risk measures, $\hat{\mathbb{D}}_i$, extracted from option prices. This factor is interpreted as p_t^* , the risk-neutral probability of a disaster. A monthly time-series of $\hat{\mathbb{D}}_i$ is constructed by taking the within-month average of daily $\hat{\mathbb{D}}_i$ for each firm i . A firm must have at least 70% of its daily observation in a month to be included. Principal component analysis is conducted on the correlation matrix of monthly $\hat{\mathbb{D}}_i$. This analysis applies to the set of firms with at least 80% of its monthly observations. To generate the time-series of the factor, I fill missing data in as their mean. The results are not sensitive to these choices. Full details of the procedure are contained in Appendix B.6. The date range is from January 1996 to May 2015 and the frequency is monthly.

Figure 3: Calibrated Disaster- β



Notes: This figure displays the model-implied disaster- β 's as a function of firm recovery in a disaster, F_i . Small values of F_i indicate firm i experiences large losses in a disaster. Large and positive values of β_i indicate that firm i 's equity returns exhibit positive comovement with changes in the risk-neutral likelihood of a disaster risk. Disaster- β 's are computed according to the expression in Corollary 1. The parameter values that generate the β_i are listed along the top of the figure.

Table 1: Correlation Analysis of p_t^*

Panel A: Correlation Analysis of Levels

	p_t^*	\mathbb{D}_{SPX}	P_t/E_t	VRP_t	TED_t	VIX_t
p_t^*	1.00	0.88	-0.23	0.29	0.50	0.85
\mathbb{D}_{SPX}		1.00	-0.24	0.40	0.56	0.84
P_t/E_t			1.00	-0.01	0.13	-0.11
VRP_t				1.00	0.30	0.51
TED_t					1.00	0.50
VIX_t						1.00

Panel B: Correlation Analysis of First Differences

	p_t^*	\mathbb{D}_{SPX}	P_t/E_t	VRP_t	TED_t	VIX_t	LEV_t
p_t^*	1.00	0.69	-0.49	-0.11	0.35	0.63	-0.26
\mathbb{D}_{SPX}		1.00	-0.51	0.19	0.39	0.79	-0.26
P_t/E_t			1.00	-0.15	-0.26	-0.75	0.11
VRP_t				1.00	0.26	0.29	0.42
TED_t					1.00	0.42	-0.12
VIX_t						1.00	-0.14
LEV_t							1.00

Notes: This table displays the correlation between p_t^* and various other financial market variables. The top panel displays correlations in levels, and the bottom panel displays correlations in first differences. \mathbb{D}_{SPX} is the disaster-risk measure defined in Equation (20), computed for the S&P 500. P_t/E_t is the cyclically adjusted price-earnings ratio (CAPE) for the S&P 500, obtained from Robert Shiller's website. When first differencing P_t/E_t , I use log-differences. VRP_t is the variance risk premium, obtained from Hao Zhou's website (series EVRP). TED_t is the TED spread and VIX_t is the VIX volatility index. LEV_t is the leverage factor of Adrian et al. (2014). In the top panel, LEV_t is excluded because it represents innovations to intermediary leverage. In addition, LEV_t is computed quarterly through 2009Q4, so all correlations with this series are computed by first averaging monthly measures within each quarter, then first differencing. All other data is monthly and spans January 1996 to May 2015.

Table 2: The Likelihood of a Disaster and Future Economic Activity

Panel A: ΔUnemployment$_{t+1}$				
Financial Indicator	Full Sample		Excluding Crisis	
	(1)	(2)	(3)	(4)
Term Spread $_t$	-0.004 (-0.27)	-0.000 (-0.02)	-0.001 (-0.07)	0.000 (0.00)
Real FFR $_t$	0.008 (0.58)	0.012 (0.74)	0.017 (1.01)	0.018 (1.12)
Stock Market Return $_t$	-0.017 (-1.87)	-0.019 (-2.11)	-0.012 (-1.19)	-0.014 (-1.33)
p_t^*	0.065 (5.04)	0.079 (4.07)	0.056 (2.02)	0.081 (2.53)
VIX $_t$		-0.017 (-0.86)		-0.018 (-0.87)
Adjusted R^2	26.0	25.9	5.7	5.6

Panel B: Δln(Industrial Production)$_{t+1}$				
Financial Indicator	Full Sample		Excluding Crisis	
	(1)	(2)	(3)	(4)
Term Spread $_t$	0.128 (2.67)	0.129 (2.45)	0.135 (2.54)	0.134 (2.48)
Real FFR $_t$	0.098 (1.92)	0.098 (1.85)	0.097 (1.71)	0.096 (1.67)
Stock Market Return $_t$	0.075 (1.84)	0.075 (1.96)	0.027 (0.84)	0.030 (0.91)
p_t^*	-0.158 (-3.80)	-0.156 (-1.61)	-0.162 (2.14)	-0.196 (-1.44)
VIX $_t$		-0.001 (-0.02)		0.022 (0.31)
Adjusted R^2	25.5	25.2	9.6	9.2

Panel C: ΔTotal Capacity Utilization$_{t+1}$				
Financial Indicator	Full Sample		Excluding Crisis	
	(1)	(2)	(3)	(4)
Term Spread $_t$	0.109 (2.53)	0.110 (2.67)	0.111 (2.82)	0.112 (2.72)
Real FFR $_t$	0.022 (0.73)	0.023 (0.64)	0.008 (0.22)	0.009 (0.23)
Stock Market Return $_t$	0.048 (1.56)	0.047 (1.55)	0.009 (0.40)	0.008 (0.30)
p_t^*	-0.114 (-4.11)	-0.108 (-1.67)	-0.093 (-1.82)	-0.073 (-0.70)
VIX $_t$		-0.008 (-0.15)		-0.014 (-0.25)
Adjusted R^2	23.0	22.6	16.2	15.8

Panel D: Δln(Manuf. Inventories)$_{t+1}$				
Financial Indicator	Full Sample		Excluding Crisis	
	(1)	(2)	(3)	(4)
Term Spread $_t$	2.141 (3.20)	1.950 (3.09)	2.186 (3.50)	2.142 (3.30)
Real FFR $_t$	1.108 (2.86)	0.878 (1.60)	0.929 (1.67)	0.869 (1.47)
Stock Market Return $_t$	-0.050 (-0.18)	0.077 (0.21)	-0.040 (-0.11)	0.080 (0.21)
p_t^*	-1.597 (8.96)	-2.475 (-3.52)	-0.611 (-0.77)	-2.081 (-1.70)
VIX $_t$		1.063 (1.41)		1.044 (1.24)
Adjusted R^2	18.6	18.9	19.8	19.9

Notes: This table reports forecasting regressions of the general form $\Delta Y_{t+1} = c + \sum_{i=1}^4 \beta_i \Delta Y_{t+1-i} + \sum_k \gamma_k X_k$, where Y_{t+1} denotes an indicator of economic activity at month $t + 1$. When forecasting the growth of industrial production and the manufacturing inventories index, Y_t is the natural logarithm of each series and ΔY_t is scaled by 100. The unemployment and total capacity utilization series are expressed in percentages. The X_k variables are as follows: $RF\text{FR}_t$ is the real federal funds rate in month t , computed as the average daily federal funds rate within month t minus the log-difference of core PCE inflation between month $t - 1$ and $t - 13$. R_t^{CRSP} is the excess return of the CRSP value-weighted index and TS_t is the difference between the 10 year and 3 month constant maturity Treasury yield, both measured at the end of month t . p_t^* is the estimated likelihood of a disaster in month t . All X variables have been standardized to have mean zero and variance one. Regressions are run at a one month frequency. The adjusted R^2 statistic is reported in percentages. Standard errors are calculated according to Newey-West (1987), with lags chosen according to the Newey-West bandwidth method and weights according to the Bartlett kernel. t -statistics for each regression are listed in below point estimates parenthesis. The full sample ranges from April 1996 to April 2015, and the sample excluding the crisis omits the period of December 2007 to June 2009.

Table 3: Summary Statistics for Equal-Weighted Portfolios Sorted on Disaster Risk Betas

	<i>Low → High Disaster Risk</i>				
	1	2	3	4	5
Annualized Excess Return (%)	5.08	8.25	10.53	11.60	12.26
Annualized Volatility (%)	29.37	19.76	17.08	17.28	22.29
Annualized Sharpe Ratio	0.17	0.42	0.62	0.67	0.55
Market Cap (\$bn)	2.25	4.29	4.50	3.73	1.47
Disaster Beta	13.63	4.77	1.11	-2.47	-11.52
Avg. Number of Firms	897	896	896	896	897

Notes: This table displays summary statistics for portfolios sorted on disaster risk betas. At the end of each month, eligible stocks are sorted into 5 equal-weighted portfolios according to their disaster risk beta. The betas are constructed in real-time, using weekly data over the previous year, via the following regression: $r_{it} = a_i + \beta_i \Delta p_t^* + e_{it}$. Disaster risk is decreasing in disaster beta, so high disaster risk stocks have the most negative disaster risk betas. Eligible stocks are defined as ordinary common shares listed on the NYSE, NYSE AMEX, or NASDAQ with a full set of weekly returns over the previous year and a share price of at least \$1. Market capitalization is reported in billions of dollars. Both market capitalization and disaster risk betas are computed as equal-weighted averages within each quintile. Return data is monthly and spans February 1997 to December 2014.

Table 4: Alphas for Equal-Weighted Portfolios Sorted on Disaster Risk Betas

	<i>Low → High Disaster Risk</i>					
	1	2	3	4	5	5-1
Annualized α (%)	-2.97	0.50	3.11	4.48	4.59	7.56
	(-1.30)	(0.33)	(2.13)	(2.80)	(1.93)	(2.35)
CAPM- β	1.09	0.89	0.79	0.78	0.88	-0.21
	(14.07)	(20.68)	(21.71)	(18.30)	(15.04)	(-1.93)
Size- β	1.03	0.65	0.56	0.55	0.82	-0.21
	(19.48)	(11.25)	(10.93)	(8.66)	(12.33)	(-2.52)
Value- β	-0.26	0.23	0.38	0.37	0.23	0.49
	(-1.62)	(3.20)	(8.30)	(6.45)	(2.76)	(2.20)
Momentum- β	-0.31	-0.15	-0.12	-0.16	-0.27	0.04
	(-3.38)	(-3.12)	(-3.90)	(-4.02)	(-4.79)	(0.31)
$H_0 : \text{All } \alpha\text{'s} = 0$	$p\text{-value} = 1.28\%$					

Notes: This table displays risk-adjusted excess returns (alphas) and risk-factor loadings for portfolios sorted on disaster risk betas. At the end of each month, eligible stocks are sorted into 5 equal-weighted portfolios according to their disaster risk beta. The betas are constructed in real-time, using weekly data over the previous year, via the following regression: $r_{it} = \alpha_i + \beta_i \Delta p_t^* + e_{it}$. Disaster risk is decreasing in disaster beta, so high disaster risk stocks have the most negative disaster risk betas. Eligible stocks are defined as ordinary common shares listed on the NYSE, NYSE AMEX, or NASDAQ with a full set of weekly returns over the previous year and a share price of at least \$1. The reported α is the intercept of a time-series regression of excess portfolio returns on the Fama-French factor returns and momentum factor returns. The reported β 's are the loadings on each respective factor. The Fama-French and momentum factors were obtained from Ken French's website. Return data is monthly and spans February 1997 to December 2014. t -statistics are listed below their respective point estimates in parenthesis and are computed according to Newey-West (1987). The reported p -value corresponds to the joint test that all $\alpha_1 = \dots = \alpha_5 = 0$.

Table 5: Risk-Adjusted Returns for Value-Weighted Portfolios Sorted on Disaster Risk Betas

	<i>Low → High Disaster Risk</i>					
	1	2	3	4	5	5-1
Annualized α (%)	-8.89	-3.39	-0.02	0.84	2.89	11.78
	(-3.73)	(-2.07)	(-0.02)	(0.58)	(1.35)	(3.11)
CAPM- β	1.27	0.97	0.82	0.81	1.01	-0.26
	(10.58)	(19.36)	(29.09)	(21.45)	(14.49)	(-1.51)
Size- β	0.20	-0.08	0.00	-0.06	0.13	-0.06
	(2.34)	(-1.26)	(0.09)	(-1.21)	(2.11)	(-0.54)
Value- β	-0.40	-0.03	0.13	0.24	0.30	0.70
	(-1.67)	(-0.33)	(3.77)	(3.37)	(2.04)	(1.86)
Momentum- β	-0.22	0.01	-0.03	-0.04	-0.09	0.14
	(-2.22)	(0.24)	(-0.94)	(-0.80)	(-1.08)	(0.78)
H_0 : All α 's = 0	p -value = 0.53%					

Notes: This table displays risk-adjusted excess returns (alphas) and risk-factor loadings for portfolios sorted on disaster risk betas. At the end of each month, eligible stocks are sorted into 5 value-weighted portfolios according to their disaster risk beta. The betas are constructed in real-time, using weekly data over the previous year, via the following regression: $r_{it} = \alpha_i + \beta_i \Delta p_t^* + e_{it}$. Disaster risk is decreasing in disaster beta, so high disaster risk stocks have the most negative disaster risk betas. Eligible stocks are defined as ordinary common shares listed on the NYSE, NYSE AMEX, or NASDAQ with a full set of daily returns over the previous quarter and a share price of at least \$1. The reported α is the intercept of a time-series regression of excess portfolio returns on the Fama-French factor returns and momentum factor returns. The reported β 's are the loadings on each respective factor. The Fama-French and momentum factors were obtained from Ken French's website. Return data is monthly and spans February 1997 to December 2014. t -statistics are listed below their respective point estimates in parenthesis and are computed according to Newey-West (1987). The reported p -value corresponds to the joint test that all $\alpha_1 = \dots = \alpha_5 = 0$.

Table 6: Alphas for Equal-Weighted Portfolios Sorted on Disaster Risk (Daily) Betas

	<i>Low → High Disaster Risk</i>					
	1	2	3	4	5	5-1
Annualized α (%)	-2.79	0.79	2.44	3.20	3.36	6.16
	(-1.23)	(0.52)	(1.74)	(2.01)	(1.55)	(2.03)
CAPM- β	1.19	0.96	0.82	0.74	0.78	-0.41
	(15.02)	(21.51)	(18.85)	(19.20)	(20.99)	(-4.51)
Size- β	1.11	0.75	0.57	0.53	0.69	-0.42
	(19.07)	(15.32)	(11.48)	(9.60)	(8.52)	(-4.11)
Value- β	-0.27	0.21	0.33	0.35	0.14	0.41
	(-1.56)	(3.14)	(6.73)	(6.64)	(1.99)	(2.00)
Momentum- β	-0.31	-0.15	-0.13	-0.13	-0.28	0.04
	(-3.16)	(-3.49)	(-3.89)	(-3.36)	(-6.27)	(0.34)
$H_0 : \text{All } \alpha\text{'s} = 0$	$p\text{-value} = 17.4\%$					

Notes: This table displays risk-adjusted excess returns (alphas) and risk-factor loadings for portfolios sorted on disaster risk betas. At the end of each month, eligible stocks are sorted into 5 equal-weighted portfolios according to their disaster risk beta. The beta for firm i is constructed in real-time, using daily data over the previous quarter, via the following regression: $r_{it} = a_i + \sum_{k=0}^2 \beta_{ik} \Delta p_{t-k}^* + e_{it}$. The disaster- β for firm i is then $\sum_{k=0}^2 \beta_{ik}$. Disaster risk is decreasing in disaster beta, so high disaster risk stocks have the most negative disaster risk betas. Eligible stocks are defined as ordinary common shares listed on the NYSE, NYSE AMEX, or NASDAQ with a full set of daily returns over the previous quarter and a share price of at least \$1. The reported α is the intercept of a time-series regression of excess portfolio returns on the Fama-French factor returns and momentum factor returns. The reported β 's are the loadings on each respective factor. The Fama-French and momentum factors were obtained from Ken French's website. Return data is monthly and spans February 1997 to December 2014. t -statistics are listed below their respective point estimates in parenthesis and are computed according to Newey-West (1987). The reported p -value corresponds to the joint test that all $\alpha_1 = \dots = \alpha_5 = 0$.

Table 7: Alphas for Value-Weighted Portfolios Sorted on Disaster Risk (Daily) Betas

	<i>Low → High Disaster Risk</i>					
	1	2	3	4	5	5-1
Annualized α (%)	-7.00	-3.14	-2.51	-0.36	3.03	10.03
	(-3.13)	(-2.07)	(-1.69)	(-0.22)	(1.52)	(3.40)
CAPM- β	1.33	1.02	0.89	0.77	0.81	-0.52
	(13.98)	(19.35)	(21.97)	(21.07)	(16.02)	(-4.08)
Size- β	0.45	0.08	-0.05	-0.11	-0.03	-0.48
	(6.05)	(1.58)	(-1.03)	(-3.21)	(-0.45)	(-4.32)
Value- β	-0.41	-0.03	0.18	0.20	0.21	0.62
	(-2.13)	(-0.60)	(4.94)	(2.91)	(2.08)	(2.17)
Momentum- β	-0.14	-0.00	0.01	0.02	-0.11	0.03
	(-1.32)	(-0.01)	(0.70)	(0.35)	(-1.58)	(0.18)
H_0 : All α 's = 0	p -value = 0.75%					

Notes: This table displays risk-adjusted excess returns (alphas) and risk-factor loadings for portfolios sorted on disaster risk betas. At the end of each month, eligible stocks are sorted into 5 value-weighted portfolios according to their disaster risk beta. The beta for firm i is constructed in real-time, using daily data over the previous quarter, via the following regression: $r_{it} = a_i + \sum_{k=0}^2 \beta_{ik} \Delta p_{t-k}^* + e_{it}$. The disaster- β for firm i is then $\sum_{k=0}^2 \beta_{ik}$. Disaster risk is decreasing in disaster beta, so high disaster risk stocks have the most negative disaster risk betas. Eligible stocks are defined as ordinary common shares listed on the NYSE, NYSE AMEX, or NASDAQ with a full set of daily returns over the previous year and a share price greater than \$1. The reported α is the intercept of a time-series regression of excess portfolio returns on the Fama-French factor returns and momentum factor returns. The reported β 's are the loadings on each respective factor. The Fama-French and momentum factors were obtained from Ken French's website. Return data is monthly and spans February 1997 to December 2014. t -statistics are listed below their respective point estimates in parenthesis and are computed according to Newey-West (1987). The reported p -value corresponds to the joint test that all $\alpha_1 = \dots = \alpha_5 = 0$.

Table 8: The Impact of Changes in Disaster Likelihood on the S&P 500

	Dependent Variable:	
	(1)	(2)
	r_t	$\Delta \log(CAPE)_t$
Δp_t^*	-6.37 (-7.91)**	-6.14 (-7.62)**
Constant	0.005 (1.71)*	0.0003 (0.10)
R^2	25.2%	23.6%
N	232	232

Notes: The table reports results from regressions of the form $y_t = c + \beta \Delta p_t^*$ for two different dependent variables. In column (1), y_t corresponds to the realized return of the S&P 500. In column (2), y_t is the growth rate of the cyclically adjusted price-earnings ratio (CAPE) for the S&P 500. All regressions are run at a one month frequency spanning the period of February 1996 to May 2015. The price and CAPE data was obtained from Robert Shiller's website. Standard errors are calculated according to Newey-West (1987), with lags chosen according to the Newey-West bandwidth method and weights according to the Bartlett kernel. t -statistics for each regression are listed in parenthesis. **, * indicates coefficient is statistically different than zero at the 5 percent and 10 percent confidence level, respectively.

Table 9: Model-Implied vs. Realized Recovery Rates for Disaster Portfolios

Portfolio	Realized Recovery (\widehat{F}_i)	Model-Implied Recovery (F_i)
1	0.46	0.76
2	0.55	0.62
3	0.58	0.51
4	0.58	0.46
5	0.49	0.43

Notes: This table displays model-implied versus realized recovery rates for portfolios sorted on disaster risk betas. Model-implied recovery rates are computed according to $F_i = 1 - XR_i / (pB^{-\gamma})$, where XR_i is the average excess return for portfolio i , p is the probability of disaster, B is the recovery rate of consumption during disasters, and γ is the risk-aversion parameter. Following Nakamura et al. (2013), I set $p = 0.028$. In addition, I set $B = 0.60$, and $\gamma = 4$. The realized recovery rates are computed as one minus the minimum 12-month cumulative return for each portfolio over the period of December 2007 to June 2009.

Table 10: Filters Applied to Individual Options

Filter	Reason	Example of Source (if possible)
Open interest is greater than zero	Avoid thinly traded options	Bakshi et al. (2003) - BKM
Bid price is greater than zero	Avoid thinly traded options	BKM and sources therein
Bid-Ask spread is positive	Avoid spurious options	Du and Kapadia (2013) - DK
Calls: $K > P_{it}$ and Puts: $K < P_{it}$	OTM options for liquidity	BKM, DK
Calls: $\mathcal{C}_i(t, \tau; K) < P_{it}$	No-arbitrage	
Calls: $\mathcal{C}_i(t, \tau; K) \geq \max(0, P_{it} - PV(K) - PV(D_i(t, \tau)))$	No-arbitrage	
Puts: $\mathcal{P}_i(t, \tau; K) < K$	No-arbitrage	
Puts: $\mathcal{P}_i(t, \tau; K) \geq \max(0, K - P_{it})$	No-arbitrage. Trivially satisfied	
Puts: $\mathcal{P}_i(t, \tau; K) \geq \max(0, PV(K) + PV(D_i(t, \tau)) - P_{it})$	No-arbitrage	
Option has standard settlement	Use standard options	
Option Implied Volatility and Delta Listed	Avoid spurious options	

Notes: This table lists all of the filters that are applied to each firm's option price in order to compute daily measures of $\hat{\mathbb{D}}_{it}$ as defined in Equation (20). Please refer to Appendix B for more details.

References

- Tobias Adrian, Erkki Etula, and Tyler Muir. Financial intermediaries and the cross-section of asset returns. *The Journal of Finance*, 69(6):2557–2596, 2014.
- David Backus, Mikhail Chernov, and Ian Martin. Disasters implied by equity index options. *The Journal of Finance*, 66(6):1969–2012, 2011.
- Gurdip Bakshi, Nikunj Kapadia, and Dilip Madan. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1):101–143, 2003.
- Robert J. Barro. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866, 2006. doi: 10.1162/qjec.121.3.823.
- Robert J. Barro. Rare disasters, asset prices, and welfare costs. *American Economic Review*, 99(1):243–64, 2009.
- Robert J. Barro and Tao Jin. On the size distribution of macroeconomic disasters. *Econometrica*, 79(5):1567–1589, 2011.
- Robert J. Barro and Jose F. Ursua. Consumption disasters in the twentieth century. *American Economic Review*, 98(2):58–63, 2008.
- Tim Bollerslev, Viktor Todorov, and Lai Xu. Tail risk premia and return predictability. *Journal of Financial Economics*, 118(1):113 – 134, 2015.
- Mark Britten-Jones and Anthony Neuberger. Option prices, implied price processes, and stochastic volatility. *The Journal of Finance*, 55(2):839–866, 2000.
- John Y. Campbell and John H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, 1999.
- Peter Carr and Dilip Madan. Towards a theory of volatility trading. *Option Pricing, Interest Rates and Risk Management, Handbooks in Mathematical Finance*, pages 458–476, 2001.
- Peter Carr and Liuren Wu. Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341, 2009.
- Hui Chen, Winston Wei Dou, and Leonid. Kogan. Measuring the “dark matter” in asset pricing models. *Working Paper*, 2015.
- John C Cox, Jonathan E Ingersoll, and Stephen A Ross. A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407, 1985.

- Martijn Cremers, Michael Halling, and David Weinbaum. Aggregate jump and volatility risk in the cross-section of stock returns. *The Journal of Finance*, 70(2):577–614, 2015.
- Patrick Dennis and Stewart Mayhew. Risk-neutral skewness: Evidence from stock options. *Journal of Financial and Quantitative Analysis*, 37(03):471–493, 2002.
- Jian Du and Nikunj Kapadia. The tail in the volatility index. *Working Paper*, 2013.
- Darrell Duffie and Larry G Epstein. Stochastic differential utility. *Econometrica*, pages 353–394, 1992.
- Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):pp. 937–969, 1989.
- Emmanuel Farhi, Samuel Paul Fraiberger, Xavier Gabaix, Romain Ranciere, and Adrien Verdelhan. Crash risk in currency markets. (15062), June 2014.
- Andrea Frazzini and Lasse Heje Pedersen. Betting against beta. *Journal of Financial Economics*, 111(1):1–25, 2014.
- Xavier Gabaix. Linearity-generating processes: A modelling tool yielding closed forms for asset prices. *NYU Working Paper*, 2009.
- Xavier Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics*, 127(2):645–700, 2012.
- George Gao and Zhaogang Song. Rare disaster concerns everywhere. *Working Paper*, 2015.
- Simon Gilchrist and Egon Zakrajsek. Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720, 2012.
- Francois Gourio. Disaster risk and business cycles. *American Economic Review*, 102(6):2734–66, September 2012.
- Christian Julliard and Anisha Ghosh. Can rare events explain the equity premium puzzle? *Review of Financial Studies*, 25(10):3037–3076, 2012.
- Bryan Kelly and Hao Jiang. Tail risk and asset prices. *Review of Financial Studies*, 2014.
- Mete Kilic and Jessica A. Wachter. Risk, unemployment, and the stock market: A rare-event-based explanation of labor market. *Working Paper*, 2015.
- Ian Martin. Simple variance swaps. *Working Paper*, 2013.

- Lior Menzly, Tano Santos, and Pietro Veronesi. Understanding predictability. *Journal of Political Economy*, 112(1):1–47, 2004.
- Emi Nakamura, Jon Steinsson, Robert Barro, and Jose Ursua. Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics*, 5(3): 35–74, 2013.
- Jun Pan. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of financial economics*, 63(1):3–50, 2002.
- Thomas A. Rietz. The equity risk premium: A solution. *Journal of Monetary Economics*, 22(1): 117–131, July 1988.
- Pedro Santa-Clara and Shu Yan. Crashes, volatility, and the equity premium: Lessons from s&p 500 options. *The Review of Economics and Statistics*, 92(2):435–451, 2010.
- Sang Byung Seo and Jessica A. Wachter. Option prices in a model with stochastic disaster risk. *Working Paper*, 2015.
- Jessica A Wachter. Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, 68(3):987–1035, 2013.

A Appendix: A Full Model of Rare Disasters

In the Section 2 of the main text, I presented the skeleton of the rare disasters model that I repeatedly use to interpret my empirical findings. In this appendix, I present a more complete version of the model, which is a continuous time generalization of the model in [Gabaix \(2012\)](#).²⁶ To make this appendix relatively self-contained, I repeat many of the primitives in the model that I outlined in Section 2 and then proceed to compute the important equilibrium quantities.

A.1 Macroeconomic Environment

I begin with an endowment economy inhabited by an infinitely lived representative agent. The endowment follows the stochastic process:²⁷

$$\frac{dC_t}{C_{t-}} = g_C dt + \sigma_C dW_{C,t} + (B_t - 1)dJ_t, \quad (4)$$

where $W_{C,t}$ is a standard Brownian motion and dJ_t is a jump process that represents macroeconomic disasters. dJ_t takes a value of one if a disaster occurs and zero otherwise. The Poisson arrival rate of disasters is time varying and denoted by p_t . Hence, a disaster happens over a small time interval, dt , with a probability $p_t dt$. The variable $B_t > 0$ captures the magnitude of a consumption disaster when it occurs. For example, a value of $B_t = 0.8$ means that when a disaster occurs, aggregate consumption drops by 20%. I further assume that the intensity process describing disaster arrivals, p_t , is independent of the magnitude of a disaster, B_t . The Brownian motion $W_{C,t}$ is independent of dJ_t and B_t , which means that in the absence of disasters (i.e. normal times) consumption growth is i.i.d log-normal with a volatility of σ_C .

I assume the representative agent has recursive preferences as in [Epstein and Zin \(1989\)](#) and [Duffie and Epstein \(1992\)](#). The main advantage of recursive preferences is the ability to decouple the agent's elasticity of intertemporal substitution, ψ , from risk aversion, γ . In a discrete time setting, Epstein and Zin (1989) show that the representative agent's stochastic discount factor can be expressed as:

$$\frac{M_{t+\tau}}{M_t} = e^{-\frac{\rho}{\chi}\tau} \left(\frac{C_{t+\tau}}{C_t} \right)^{-\frac{1}{\chi\psi}} R_{C,t+\tau}^{1/\chi-1} \quad (5)$$

where $\chi = (1 - 1/\psi)/(1 - \gamma)$ and $R_{C,t+\tau}$ is the gross return of a claim to aggregate consumption between t and $t + \tau$. Here, ρ captures the rate of time preference of the agent. When solving the model, I obtain the stochastic discount factor by taking the continuous time limit of Equation (5), i.e. as $\tau \searrow 0$.

²⁶Like [Gabaix \(2012\)](#), I take repeated advantage of the linearity-generating (LG) class of stochastic process pioneered by [Gabaix \(2009\)](#), which the reader should refer to for a more detailed exposition of LG processes.

²⁷To be precise, C_{t-} denotes $\lim_{s \uparrow t} C_s$.

A.2 Setup for Stocks

A generic stock i is defined as a claim to a stream of dividends $(D_{it})_{t \geq 0}$, where dividends grow according to:

$$\frac{dD_{it}}{D_{it-}} = g_{iD}dt + dN_{it}^D + (F_{it} - 1)dJ_t \quad (6)$$

The dividend growth shock N_{it}^D is a martingale so that $\mathbb{E}_t[dN_{it}^D] = 0$ and is independent of the disaster event. The higher moments of N_{it}^D are left unspecified for now. As we will see shortly, only the first moment of dN_{it}^D is important for determining the *level* of asset prices. In normal times, dividends grow at an expected rate of g_{iD} , but if there is a disaster, the stock dividend experiences an additional shock, F_{it} . $F_{it} > 0$ captures how much dividends recover during a disaster. For instance, if $F_{it} = 0.9$, then dividends drop by 10% when a disaster arrives. Firms with high values of F_{it} therefore perform better during a crisis.

In order to compute equilibrium, we must put more structure on the arrival rate of disasters (p_t), the recovery rate of disasters at the macroeconomic level (B_t), and the dividend recovery rate during a disaster (F_{it}). One way to accomplish this is to model the dynamics of p_t , B_t , and F_{it} separately. Wachter (2013) is one example of this approach, as she models p_t as a Cox, Ingersoll, and Ross (1985) process. Alternatively, Gabaix (2012) models the dynamics of p_t , B_t , and F_{it} jointly, which I adopt in continuous time here. I do so through a variable called resilience, H_{it} , which summarizes how a firm is expected to perform if a disaster occurs over the next instant, dt . I define stock i 's resilience exactly as in Equation (31) in Appendix C; however, in Appendix C I also show that a very good approximation that is more useful for understanding the underlying economics in the model is:

$$H_{it} = p_t \mathbb{E}_{t-} [B_t^{-\gamma} (F_{it} + (\chi - 1)B_t) - \chi] + \frac{\sigma_C^2}{2} \gamma(1 + \gamma) \quad (7)$$

It is easy to see that firms that are expected to perform well in a disaster (i.e. firms whose dividends recover a lot) have high values of F_{it} , and thus are high resilience stocks. For ease of notation, I drop the limit notation of \mathbb{E}_{t-} and replace it by \mathbb{E}_t , except when it is necessary to distinguish the two.

Resilience itself is modeled as a linearity generating (LG) process as in Gabaix (2009). Linearity generating processes behave, in an approximate sense, like an AR(1) process. The distinction between an LG process and an AR(1) process is subtle; however, LG processes are useful in economics because they deliver linear expressions for equilibrium price-dividend ratios without needing log-linearizations. As is common in macroeconomics, I further decompose resilience into a long-run mean, H_{i*} , and deviations from that mean, \hat{H}_{it} :

$$H_{it} = H_{i*} + \hat{H}_{it} \quad (8)$$

In turn, firm i 's resilience (deviations from mean) evolves according to the following stochastic differen-

tial equation:²⁸

$$\begin{aligned} d\hat{H}_{it} &= -\phi_{H,i}\hat{H}_{it}dt - \hat{H}_{it}^2 dt + dN_{it}^H \\ &\approx -\phi_{H,i}\hat{H}_{it}dt + dN_{it}^H \end{aligned} \quad (9)$$

where N_{it}^H is a martingale. It is easy to now see the relationship between an AR(1) process and an LG process. The term $-\phi_{H,i}\hat{H}_{it}dt$ would be present in the continuous time analogue of an AR(1) process, and so $\phi_{H,i}$ determines the speed of mean reversion for the resiliency process. It is sensible that resilience should behave similarly to an AR(1) process — if a firm has high resilience today, it is likely that resilience tomorrow is also high. The “extra” term, $-\hat{H}_{it}^2 dt$, is unique to the LG class, and in practice is quite small because \hat{H}_{it} captures deviations from the mean of resilience. As noted in [Gabaix \(2009\)](#), there are restrictions on the variance of the noise process dN_{it}^H that ensure prices remain positive. In [Appendix C.3](#), I provide one example of a process that satisfies the restrictions and is also internally consistent with the rest of the model. With the primitives of the model now specified, I now turn to analyzing equilibrium asset prices in this economy.

A.3 Equilibrium Stock Prices and Returns

Theorem 1 solves for equilibrium price-dividend ratios.

Theorem 1. *Define $\delta = \rho + g_C/\psi$ and $\delta_i = \delta - g_{iD} - H_{i^*}$. δ_i can be interpreted as a stock’s effective discount rate. Then in equilibrium, the price-dividend ratio of stock i is:*

$$\frac{P_{it}}{D_{it}} = \frac{1}{\delta_i} \left(1 + \frac{\hat{H}_{it}}{\delta_i + \phi_{H,i}} \right) \quad (10)$$

Proof. The proof can be found in [Appendix C](#). □

The equilibrium price-dividend ratio given by Equation (10) has a very intuitive form. First consider what determines the stock’s effective discount rate, δ_i . All else equal, firms with a high level of average resilience, H_{i^*} , have a low discount rate and therefore a higher valuation ratio. This result makes sense as these are the firms whose dividend growth is expected to recover more during a disaster. Similarly, when a firm’s resilience is above average (i.e. $\hat{H}_{it} > 0$), its valuation ratio will be above average. Since resilience mean reverts, this generates mean-reversion in the price-dividend ratio and hence predictability in stock prices. Naturally, a stock with a high level of resilience has a low equity premium (and a high price-dividend ratio); thus, high valuation ratios forecast lower subsequent returns, as we see in the data.

A natural quantity of interest that derives from the equilibrium price dividend ratio in Equation (10) is the equilibrium return process for firm i ’s stock. This can be computed under two different probability measures. The first is the physical measure, which corresponds to the actual distribution of return realizations. The second is the so-called risk-neutral measure, which weights states of nature according to marginal utility in those states. Naturally, states in which marginal utility is high (i.e. disaster states) have

²⁸To be precise, the following specification ignores covariance of the SDF with resilience shocks in normal times. See [Section C](#) for the version of the model where this is not the case.

a higher risk-neutral likelihood than a physical likelihood. Intuitively, risk-neutral probabilities are just Arrow-Debreu state prices, appropriately normalized so that they correspond to a probability measure. The risk-neutral likelihood will prove useful when examining the structure of option prices implied by the disaster economy. The following proposition characterizes firm-level returns under both measures:

Theorem 2. *Given the equilibrium price-dividend ratio in Equation (10), the process for dividends in Equation (6), and the process for resilience in Equation (9), equilibrium returns for firm i evolve as follows:*

$$\frac{dP_{it}}{P_{it^-}} = \mu_{it}dt + \sigma_{it}dW_{it} + (F_{it} - 1)dJ_t \quad (11)$$

where μ_{it} and σ_{it} are defined in Equation (35), and dW_{it} is a Brownian motion.

Next, let the superscript $*$ denote a stochastic process under the risk-neutral measure \mathbb{Q} . The measure \mathbb{Q} is defined formally by Equation (41). Then the risk-neutral return process evolves according to:

$$\frac{dP_{it}^*}{P_{it}^*} = r_{ft}dt + \sigma_{it}dW_{it}^* - p_t\mathbb{E}_t[B_t^{-\gamma}(F_{it} - 1)]dt + (F_{it} - 1)dJ_t^* \quad (12)$$

where r_{ft} is the risk-free short rate process given by Equation (38). The process dW_{it}^* is a Brownian motion under the risk-neutral measure. dJ_t^* is the disaster jump process under \mathbb{Q} , and has a \mathbb{Q} -arrival rate of $p_t\mathbb{E}_t[B_t^{-\gamma}]$. Furthermore, the expected excess return of the firm is, to a leading order, given by:

$$\mu_{it} - r_{ft} = p_t\mathbb{E}_t[B_t^{-\gamma}(1 - F_{it})] \quad (13)$$

Proof. The proof can be found in Appendix C.3 and C.4.²⁹ □

The first thing to note from Theorem 2 is the likelihood of a disaster under both probability measures. When disasters are negative shocks to the economy ($\mathbb{E}_t[B_t] < 0$), then the risk-neutral probability of a disaster is greater than the actual likelihood of a disaster, i.e. $p_t\mathbb{E}_t[B_t^{-\gamma}] > p_t$. This occurs because the risk-averse agent prices assets as if disasters are more likely than they actually are in nature. In addition, Theorem 2 states that the risk premium of firm i is decreasing in F_{it} . Intuitively, firms that are expected to do poorly in a disaster require higher expected returns in equilibrium. This is a testable hypothesis that I explore in Section 4.

A.4 The Risk-Neutral Disaster Intensity

The jump component in the risk-neutral return process is itself the product of a firm specific variable, $(1 - F_{it})$, and an economy wide variable, dJ_t^* . One of the primary contributions of this paper is to estimate time-series variation in the arrival rate of dJ_t^* using options data. My strategy in the main text uses more sophisticated option pricing theory, however I begin with a simple example to develop the basic intuition of my approach.

²⁹Implicit in this proof is the assumption that the shocks dN_{it}^H and dN_{it}^D are themselves governed by Brownian motions. Theorem 1 requires dN_{it}^H and dN_{it}^D to be mean-zero shocks, which holds true under this assumption. See Appendix C.3 for a specific example where all of these conditions are met.

A Simplified Example

Consider the simplest case where the volatility of firm returns σ_i , the arrival rate of disasters p , the severity of disasters B , and firm performance in a disaster F_i are all constants. In addition, suppose that we write a call and put option on firm i with a maturity of Δ , where Δ is a small unit of time. The following proposition provides a simple expression for option prices in this stylized example.

Proposition 2. *Suppose that $\sigma_{it} = \sigma_i$, $p_t = p$, $B_t = B$, and $F_{it} = F_i$ are all constants. For a small time interval, Δ , define $p^* \equiv pB^{-\gamma}\Delta$. Then options written on a stock with spot price P_{it} , strike K , volatility σ_i , and maturity Δ are approximately given by:*

$$\begin{aligned}\mathcal{P}_i(P_{it}, K) &= \theta_i \times [p^* \times \mathcal{P}_i^{BS}(F_i P_{it}, K) + (1 - p^*) \times \mathcal{P}_i^{BS}(P_{it}, K)] \\ \mathcal{C}_i(P_{it}, K) &= \theta_i \times [p^* \times \mathcal{C}_i^{BS}(F_i P_{it}, K) + (1 - p^*) \times \mathcal{C}_i^{BS}(P_{it}, K)]\end{aligned}\quad (14)$$

where $\mathcal{P}_i^{BS}(X, K)$ is the Black-Scholes value of a put option with underlying price X , strike K , volatility σ_i , maturity Δ , and interest rate 0. θ_i is a firm level constant, $\theta_i = e^{-p^*(1-F_i)}$.

Proof. See Appendix C.5. □

Under these simplifying assumptions, the value of a put (or call) is the probability weighted average of two components: (i) a Black-Scholes option with an initial price $F_i P_{it}$ and (ii) a Black-Scholes option with an initial price P_{it} . Intuitively, when there is no disaster, stock prices are log-normal and hence take their usual Black-Scholes form. When there is a disaster, stock prices are still log-normal, but the initial price is multiplied by F_i because this is the stock price after the disaster occurs. The contribution of each component to the total put price depends on the risk-neutral likelihood of a disaster.

Now assume that the firm is expected to fall in value in a disaster so that $F_i < 1$. Next, consider the following option portfolio: (i) purchase 1 unit of an out-of-the-money put option with moneyness $m_{Put} \equiv K_{Put}/P_{it} < 1$; and (ii) sell m_{Put} units of an out-of-the-money call option with moneyness $m_{Call} = m_{Put}^{-1} \equiv K_{Call}/P_{it} > 1$. For instance, if $m_{Put} = 0.9$ and $P_{it} = 100$, then the portfolio purchases one unit of a put option at strike 90, and sell 0.9 units of a call option at strike $100 \times (0.9)^{-1} = 111$. This particular portfolio is special because it isolates the disaster component of the put option. The argument proceeds in two steps.

First, observe that a good approximation of the call price is $\mathcal{C}_i(P_{it}, K_{Call}) \approx \theta_i(1 - p^*)\mathcal{C}_i^{BS}(P_{it}, K)$. The approximation is valid because the call will have virtually no value in the disaster state; the option effectively expires out of the money since the stock price drops during a disaster. In turn, all of the call option value comes from “normal-times” risk.

Second, it is easy to show that for K_{Call} and K_{Put} defined above, the following relation holds for a Black-Scholes put and call with a zero interest rate: $\mathcal{P}_i^{BS}(P_{it}, K_{Put}) = m_{Put} \times \mathcal{C}_i^{BS}(P_{it}, K_{Call})$. Intuitively in the Black-Scholes world with a zero interest rate, future risk-neutral returns are symmetric around zero. Thus, it is possible to find a combination of the Black-Scholes put and the call that exactly offsets in value.

These two observations imply that the value of the option portfolio isolates the disaster component

of the put option:

$$\begin{aligned}\mathcal{P}_i(P_{it}, K_{Put}) - m_{Put} \times \mathcal{C}_i(P_{it}, K_{Call}) &= \theta_i p^* \mathcal{P}_i^{BS}(F_i P_{it}, K_{Put}) \\ K_{Put} &= m_{Put} P_{it}, \quad m_{Put} < 1 \\ K_{Call} &= m_{Put}^{-1} P_{it}\end{aligned}$$

In practice, this portfolio resembles what is typically called a “risk-reversal”. The call option in the portfolio exactly offsets the value of the put that comes from normal-times risk, leaving only the disaster component of the put option. This stylized example is designed to provide some economic intuition as to why it is possible to use a portfolio of options to learn about the disaster states of the world. I turn now to showing that this is possible when risk-neutral returns follow more a general process, such as the one implied by the model in Equation (12).

The More Complex Version

To develop my full estimation strategy, first define the risk-neutral expected quadratic variation of the log firm stock process from the period $[t, t + \tau]$ as:

$$\mathbb{Q}\mathbb{V}_{it}(\tau) \equiv \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^{t+\tau} \sigma_{is}^2 ds + \int_t^{t+\tau} \left(\int_{\mathbb{R}^+} \ln^2(v) f_i(v; t) dv \right) p_s \mathbb{E}_t [B_s^{-\gamma}] ds \right]$$

where $f_i(F; t)$ is the probability distribution function of F_i at time t . Simply put, quadratic variation captures *any* expected squared price movements. This includes price movements that come from the diffusive portion of returns (normal-times risk) and the jump portion of returns (disaster risk). [Bakshi, Kapadia, and Madan \(2003\)](#) show how to use observed option prices to form an estimate of $\mathbb{Q}\mathbb{V}$. The details of this procedure are contained in Section [B.1](#).

Of course, there are many other measures of price variability that can be estimated from option prices. Perhaps the most popular is the so-called VIX index, which is formally defined as:

$$\mathbb{V}_{it}(\tau) \equiv \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^{t+\tau} \frac{dP_{it}}{P_{it}} - \ln \left(\frac{P_{i,t+\tau}}{P_{it}} \right) \right] \quad (15)$$

At first glance, this definition of the VIX may seem a bit obscure. Indeed, as shown by [Britten-Jones and Neuberger \(2000\)](#), when the underlying return process does not have discontinuities, $\mathbb{V}_{it}(\tau) = \mathbb{Q}\mathbb{V}_{it}(\tau)$. It is also easy to see from its definition that the absence of jumps implies quadratic variation — and hence the VIX — measures $\mathbb{E}_t^{\mathbb{Q}} \left[\int_t^{t+\tau} \sigma_{is}^2 ds \right]$. This argument underlies the usual interpretation of the VIX as a measure of variability.

Nonetheless, it is well-known that the VIX is a biased estimate of quadratic variation when the return process does contain jumps. In fact, [Carr and Wu \(2009\)](#) and [Du and Kapadia \(2013\)](#) show how one can use this bias to isolate the jump component of returns. Theorem [3](#) collects the theory behind these results more precisely (Section [B.1](#) contains the empirical implementation).

Theorem 3. Let $\mathbb{Q}\mathbb{V}_{i_t}(\tau)$ be the quadratic variation, and $\mathbb{V}_{i_t}(\tau)$ be defined as in Equation (15). Then, inside the rare disasters model, the following holds true:

$$\begin{aligned}\mathbb{D}_{i_t}(\tau) &\equiv \mathbb{Q}\mathbb{V}_{i_t}(\tau) - \mathbb{V}_{i_t}(\tau) \\ &= \Psi(f_i(F; t)) \times p_t^*(\tau)\end{aligned}\tag{16}$$

where $\Psi(f_i(F; t))$ is given by Equation (47) in Appendix C.6, and $p_t^*(\tau) = \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^{t+\tau} p_s B_s^{-\gamma} ds \right]$. $f_i(F; t)$ is the probability distribution function of the firm jump size, F_{i_t} , at time t . Furthermore, when F_{i_t} is constant, then the expression reduces to:

$$\mathbb{D}_{i_t}(\tau) = 2 \times \{1 + \ln(F_i) + [\ln(F_i)]^2 / 2 - F_i\} \times p_t^*(\tau)\tag{17}$$

Proof. The proof can be found in Appendix C.6. □

When it is clear, I suppress the arguments $\mathbb{D}_{i_t}(\tau)$ and denote it by \mathbb{D}_{i_t} for notational convenience. Similarly, I denote $p_t^*(\tau)$ by p_t^* and when the distribution of F_{i_t} is time-invariant I use Ψ_i to simplify the expression for $\Psi(f_i(F; t))$.

Theorem 3 says that all cross-sectional variation in \mathbb{D}_{i_t} is generated by cross-sectional heterogeneity in the distribution of F_{i_t} . The reason is that p_t^* is common across all firms in the economy. Furthermore, when the distribution of F_{i_t} is time-invariant — F_{i_t} can still be stochastic — then time-series variation in \mathbb{D}_{i_t} for all firms will be driven by a common factor p_t^* . This last observation is the core of my empirical strategy.

B Appendix: Data Details and Variable Construction

This appendix provides a detailed description of how I construct all the data used in the main text.

B.1 Estimating \mathbb{D}_{it}

Equation (16) defines $\mathbb{D}_{it}(\tau)$ as the difference between (risk-neutral) expected quadratic variation and the VIX. Following [Du and Kapadia \(2013\)](#), I estimate both quadratic variation and the VIX using observed option prices.³⁰

To start, [Bakshi, Kapadia, and Madan \(2003\)](#) show that a good approximation of expected quadratic variation can be formed using the observed prices of out of the money puts and calls for firm i .³¹

$$\begin{aligned} \widehat{\mathbb{QV}}_{it}(\tau) &= \frac{e^{r_{ft}\tau}}{\tau} \left[\int_{K>P_{i,t}} \frac{2(1-\ln(K/P_{i,t}))}{K^2} \mathcal{C}_{it}(P_{i,t}; K, \tau) dK \right] \\ &\quad + \frac{e^{r_{ft}\tau}}{\tau} \left[\int_{K<P_{i,t}} \frac{2(1-\ln(K/P_{i,t}))}{K^2} \mathcal{P}_{it}(P_{i,t}; K, \tau) dK \right] \\ &\quad - m_{i,t,\tau}^2 \end{aligned} \quad (18)$$

where $\mathcal{C}_{it}(P_{i,t}; K, \tau)$ is the price of a call option written on firm i , on date t , at strike K , and with a time to maturity of τ . $\mathcal{P}_{it}(P_{i,t}; K, \tau)$ defines the analogous put price. $P_{i,t}$ is the current spot price of the firm. $m_{i,t,\tau}$ is generally a small adjustment term and is defined fully in Equation (39) of [Bakshi, Kapadia, and Madan \(2003\)](#). Finally r_{ft} is the annualized risk-free rate over the life of the option. The $\hat{\cdot}$ superscript is used to denote empirical estimates from their theoretical counterparts.

In parallel, I use what is by now a standard formula to estimate the VIX.³²

$$\begin{aligned} \widehat{\text{VIX}}_{it}(\tau) &= \frac{2e^{r_{ft}\tau}}{\tau} \left[\int_{K>P_{i,t}} \frac{1}{K^2} \mathcal{C}_{it}(P_{i,t}; K, \tau) dK + \int_{K<P_{i,t}} \frac{1}{K^2} \mathcal{P}_{it}(P_{i,t}; K, \tau) dK \right] \\ &\quad - \frac{2e^{r_{ft}\tau}}{\tau} [e^{-r_{ft}\tau} (e^{r_{ft}\tau} - 1 - r_{ft}\tau)] \end{aligned} \quad (19)$$

Finally, combining Equations (18) and (19) yields the desired estimator for $\mathbb{D}_i(P_i; t, t + \tau)$:

$$\widehat{\mathbb{D}}_{it}(\tau) = \widehat{\mathbb{QV}}_{it}(\tau) - \widehat{\text{VIX}}_{it}(\tau) \quad (20)$$

³⁰There is a large options literature on estimating risk-neutral expected quadratic variance and integrated variance under a variety of return specifications. A non-exhaustive list includes [Bakshi, Kapadia, and Madan \(2003\)](#), [Carr and Madan \(2001\)](#), [Pan \(2002\)](#), and [Martin \(2013\)](#).

³¹The approximation error between Equation (18) and actual quadratic variation is related to the stochasticity of the drift of the log-return. This error goes to zero as time shrinks to zero, which is why I will consider relatively short dated options (30 day) in my empirical work.

³²The formula I present comes from [Du and Kapadia \(2013\)](#).

where I am now using the $\hat{\cdot}$ to emphasize estimates from data.

In terms of quantities within the model, the results from [Du and Kapadia \(2013\)](#) imply that $\hat{\mathbb{D}}_{it}(\tau)$ actually corresponds to an estimate for:³³

$$\hat{\mathbb{D}}_{it}(\tau) = \frac{\Psi_i}{\tau} \times p_t^*(\tau) \quad (21)$$

where I have assumed again that the distribution for F_{it} is time-invariant. Note that the only difference between Equation (21) and Equation (2) is the annualizing constant τ , which does not impact the factor analysis.

To actually construct each firm's $\mathbb{D}_{it}(\tau)$, I proceed as follows. First, for each maturity that is available, I compute the associated risk-free rate via linear interpolation over the zero-coupon yield curve provided OptionMetrics. For each option maturity τ , I require at least two OTM put and two OTM call options, which ensures some degree of symmetry when integrating over option prices.³⁴ Next, I use numerical integration to evaluate Equations (18) and (19), which I then combine to compute $\mathbb{D}_{it}(\tau)$ from Equation (20). This procedure is executed for each available maturity. Finally, I compute a $\mathbb{D}_{it}(\tau = 30 \text{ days})$ by linearly interpolating between the $\mathbb{D}_{it}(\tau)$ for maturities that straddle 30 days.

B.2 Constructing Daily $\hat{\mathbb{D}}_{it}$ for Each Firm

In order to construct the option based measures used in the paper, I first downloaded the entire OptionMetrics database at a daily frequency from 1/4/1996 to 5/28/2015. I also use the underlying security information and the risk-free rate curves from this same source.

Dividends When I consider options written on individual equity options, they are all American options. As such, there are a number of standard no-arbitrage restrictions on transaction prices that are particular to American options. Most of these include some discounting of the dividend yield. One way to throw out “spurious” option prices that violate these arbitrage restrictions is to use the dividend yields provided by OptionMetrics and then check the bounds exactly. As I argue below, in practice accounting for the effect of dividends on the early exercise feature of options is a nonissue. All of the arguments regarding dividend considerations apply to the index options in my sample as well.

Risk-Free Rates For a given date, I use the risk-free rate curve provided by OptionMetrics. This has the continuously compounded zero-coupon risk-free yield at discretely placed terms. To compute the continuously compounded riskless rate for a given term, I use: (i) linear interpolation if the terms falls within the range of terms provided by OptionMetrics or (ii) I use the rate associated with nearest provided term for terms outside of the minimum or maximum term provided. For instance, if the lowest

³³To be precise, this particular estimate ignores the contribution of the stochastic drift term. As shown in [Du and Kapadia \(2013\)](#) the “extra” terms are largely a function of the squared expected stochastic drift. For the short maturity options I consider in this paper, they then show via simulations that ignoring the stochasticity of the drift is innocuous. See Section 3.4 in [Du and Kapadia \(2013\)](#) for more details.

³⁴Additional option pricing filters can be found in Appendix B.

term provided is for 15 days, I use the 15-day rate for all terms between 1-14 days. If the riskless rates on a given date are not available, I use the most recently available set of discount rates.

B.3 First Pass Filter of the Data

I retain only options that meet all of the following criteria listed in Table 10. In addition, I list a succinct reason for each filter, as well as examples in the literature that apply the same types of filter. The operator $PV(\cdot)$ denotes the present value of its argument. $\mathcal{C}_i(t, \tau; K)$ and $\mathcal{P}_i(t, \tau; K)$ denote the price, at time t , of a call and put that expire at time $t + \tau$, with strike K , respectively. The subscript i refers to firm i . $D_i(t, \tau)$ is the dividend, if any, of the firm between date t and date $t + \tau$. The price I use for both calls and puts is the mid-point of the bid and the ask price. P_{it} is the price of the underlying firm (or index) at time t .

A few of the filters are worth further discussion:

- Satisfying the condition that $\mathcal{C}_i(t, \tau; K) \geq \max(0, P_{it} - PV(K) - PV(D_i(t, \tau)))$ is trivial. I only consider OTM call options. For reasonable discount rates, then $P_t - PV(K)$ will already be negative and the relationship is already satisfied. Subtracting out the present value of future dividends doesn't change this fact.
- The constraint that $\mathcal{P}_i(t, \tau; K) \geq \max(0, PV(K) + PV(D_i(t, \tau)) - P_{it})$ is also rather harmless. For out of the money put options, and short enough maturity puts, $PV(K) - P_t \approx K - P_{it} < 0$. Unless the dividend over the life of the option is very large, then $\max(0, PV(K) + PV(D_i(t, \tau)) - P_{it}) = 0$. In this case, discarding put options with non-zero mid-quotes is enough to satisfy the condition. Even with a dividend payment, it is unlikely it is so large so as to affect the results materially. Thus, I can safely ignore this restriction, particularly since my use of options is not for a directly tradable strategy.

After eliminating options on an individual basis, I then apply one additional filter based on all options for a firm and date. Namely, I require that there are at least two OTM put and two OTM call for each firm. Previous research using options data, e.g. [Dennis and Mayhew \(2002\)](#), show that more symmetry in terms of the number of observed calls/puts delivers better properties of implied volatility surfaces.

B.4 Building Measures of \widehat{QV} and \widehat{V}

In order to build my estimator of disaster risk, I need to build option-based measures of quadratic variation and integrated variance. In theory, these measures require observations over an infinite continuum of strikes (in order to integrate over \mathbb{R}^+). In practice this is of course not feasible. Following what has become standard practice, I adopt the same methodology that is used to construct the VIX, but I modify it to compute other higher moment estimators that I am interested in.³⁵

I use only observed option prices to compute the integrals in Equations (18) and (19). All integrals are computed numerically using trapezoidal integration. For American options, [Bakshi, Kapadia, and](#)

³⁵The exact methodology can be found at <https://www.cboe.com/micro/vix/vixwhite.pdf>. To incorporate different moment estimators, I simply alter their Equation (1) to have the appropriate formula. The only modification I make to their formula is with regard to the forward price used in formulas.

Madan (2003) show that ignoring the early-exercise premium is innocuous for short maturity options. Like with the construction of the CBOE's VIX, I begin by computing my metrics for two different maturity options. Specifically, I find the set of options that is nearest, but less than, 30 days to maturity (i.e. $\tau_1 < 30$). For these maturity options, I compute my higher moment estimators (e.g. $\widehat{\mathbb{QV}}_{it}(\tau_1)$, etc.).

Next, I find the set of options that is nearest, but greater than, 30 day days to maturity (i.e. $\tau_2 > 30$). For these maturity options, I compute my higher moment estimators (e.g. $\widehat{\mathbb{QV}}_{it}(\tau_2)$, etc.).

Finally, I compute an estimate of the "30 day" implied estimator (e.g. $\widehat{\mathbb{QV}}_{it}(\tau = 30 \text{ days})$) via linear interpolation.³⁶ If there is not a set of options with expiration below or above 30 days, I treat the available estimator as the 30 day estimate. For instance, if there are only 44 day options available (but not 16 day options), then I use the estimate coming from the 44 day options. I do not include any options whose maturity is below 7 days. The end result is a panel of 30-day implied $\widehat{\mathbb{D}}_{it}$ for 6,762,860 firm-day pairs.

B.5 Aggregating Daily $\widehat{\mathbb{D}}_{it}$ to Monthly $\widehat{\mathbb{D}}_{it}$

To smooth possible measurement errors in daily option prices, I create a monthly $\widehat{\mathbb{D}}_{it}$ for each firm based on their daily measure. Specifically, for a given month and firm i , I take a within-month average of firm i 's daily $\widehat{\mathbb{D}}_{it}$. To be included in the analysis, I require that each firm has at least 70% of its daily observations available for that month. This monthly panel of $\widehat{\mathbb{D}}_{it}$ is what I use to conduct principal component analysis and extract my monthly measure of p_t^* , which I detail next.

B.6 Constructing an Estimate of p_t^* Using the Full Dataset

The next task is to conduct principal component analysis (PCA) on the monthly panel of $\widehat{\mathbb{D}}_{it}$. To do so, I use standard PCA analysis on the correlation matrix of $\widehat{\mathbb{D}}_{it}$. Specifically, the first principal component is first eigenvector \mathbf{v} of the correlation matrix.

The usual extraction method for PCA is to multiply the $T \times N$ panel of standardized $\widehat{\mathbb{D}}_{it}$ by \mathbf{v} . By standardized, I mean that each firm's $\widehat{\mathbb{D}}_{it}$ is adjusted to have a mean zero and variance of one. I use a slightly modified procedure to extract the factor. Let me first describe it, then provide the rationale for this alternative.

I use the eigenvector \mathbf{v} to define a set of weights, $\boldsymbol{\omega}$. Thus, ω_i for firm i is simply $v_i / \sum v_i$, where v_i is the i th element of \mathbf{v} . I then compute the factor as:

$$\widehat{p}_t^* = \theta \widehat{\mathbb{D}} \boldsymbol{\omega}$$

where $\widehat{\mathbb{D}}$ is the $N \times T$ panel of $\widehat{\mathbb{D}}_{it}$. $\theta = \text{sign}(\text{correlation}(\widehat{\mathbb{D}} \boldsymbol{\omega}, \bar{\mathbb{D}}_t))$. Here, $\bar{\mathbb{D}}_t$ is the average $\widehat{\mathbb{D}}_{it}$ across all firms at each point in time.

The rationale for this construction is to maintain the units of p_t^* . This is useful when mapping a calibrated version of the model to the data. In addition, because eigenvectors are identified up scalar, multiplying by θ simply ensures that the extracted factor is positively (not negatively) correlated with the average disaster risk in the economy. This is a sensible assumption. Constructing \widehat{p}_t^* in this way does not materially impact the results, as the factor created from standard PCA is over 95% correlated with this

³⁶See page 5 in the CBOE's VIX [White Paper](#).

\hat{p}_t^* . This is also not surprising given that I am using the “usual” principal component factor loadings to instead create a weighted average.

In the model, this means that p_t^* is just a weighted average of firm-level disaster risk (consistent with the single factor structure of $(\hat{\mathbb{D}}_{it})$):

$$\begin{aligned}\hat{p}_t^* &= \sum_i \omega_i \Psi_i p_t^* \\ &= \bar{\Psi} p_t^*\end{aligned}$$

where \hat{p}_t^* is the estimate from the data and p_t^* is the model counterpart. It is easy to see here why taking a weighted-average of firm-level disaster risk is appealing: the resulting estimate is p_t^* times the economy wide average Ψ_i .

B.7 Constructing a Rolling Estimate of p_t^*

When forming portfolios that are exposed to disaster risk, it is important to make sure the strategies are tradable in real-time. In turn, this requires me to compute a real-time measure of p_t^* . I call this real-time variable a rolling measure of p_t^* , and at each point in time I use the previous year’s worth of data to construct it. In addition, I construct a rolling p_t^* at two frequencies: (i) weekly and (ii) daily. I use two frequencies (and not the monthly measure described before) so that I have a reasonable amount of observations over which to estimate disaster- β ’s.

B.7.1 Weekly Frequency

My weekly rolling estimate of p_t^* derives from the daily panel of $\hat{\mathbb{D}}_{it}$. To fix ideas, consider a particular date (daily) t for which I want to construct a rolling measure of p_t^* . I then proceed as follows:

1. Create a subsample of daily $\hat{\mathbb{D}}_{it}$ from date t to $t - 365$.
2. Eliminate firms who have less than 95% of their daily observations over the past year.
3. For the remaining firms, take a within week-firm average to compute a weekly measure of $\hat{\mathbb{D}}_{it}$ over the past year. Drop any firms who have a missing week (i.e. did not have at least one observation in a given week).
4. Conduct principal component analysis on this panel of weekly $\hat{\mathbb{D}}_{it}$. Define the rolling estimate of p_t^* using the same procedure as in Section B.6, but for this subsample.

I call this estimate $p_{RW,t}^*$, where the subscript RW indicates it is constructed on a rolling basis and at a weekly frequency.

B.7.2 Daily Frequency

My daily rolling estimate of p_t^* derives from the daily panel of $\hat{\mathbb{D}}_{it}$. To fix ideas, consider a particular date (daily) t for which I want to construct a rolling measure of p_t^* . I then proceed as follows:

1. Create a subsample of daily $\widehat{\mathbb{D}}_{it}$ from date t to $t - 365$.
2. Eliminate firms who have less than 95% of their daily observations over the past year.
3. For the remaining firms, fill any missing data with data from the previous day. If this is not possible, then leave the observation blank. Drop any firms who still have any missing data.
4. Conduct principal component analysis on this panel of daily $\widehat{\mathbb{D}}_{it}$. Define the rolling estimate of p_t^* using the same procedure as in Section B.6, but for this subsample.

I call this estimate $p_{RD,t}^*$, where the subscript RD indicates it is constructed on a rolling basis and at a daily frequency.

B.8 Forming Equity Portfolios Based on Disaster- β 's

Once I have real-time estimates of $p_{RW,t}^*$ and $p_{RD,t}^*$, I can now compute real-time disaster- β 's. The disaster- β 's at the end of each month are computed using either the previous year's weekly returns or the previous quarter's daily returns.

B.8.1 Disaster- β Estimated Using Weekly Data

Suppose we are standing on date t , which is also the end of the month. The goal is to compute disaster- β 's on which we can form portfolios to hold for the next month. For each firm i , I compute a disaster- β from weekly data using the following expression:

$$\hat{\beta}_i = \frac{\rho(r_t^i, \Delta p_{RW,t}^*) \times \sigma(r_t^i)}{\sigma(\Delta p_{RW,t}^*)}$$

where $\rho(x_t, y_t)$ is the correlation between x and y using weekly data from t to $t - 365$. $\sigma(x)$ is the volatility of x using data over this same horizon. This procedure delivers me a real-time estimate of $\hat{\beta}_i$ for each firm and at the end of each month. I compute disaster- β 's using weekly data only for firms with a full set of weekly returns over the previous year.

B.8.2 Disaster- β Estimated Using Daily Data

Suppose we are standing on date t , which is also the end of the month. The goal is to compute disaster- β 's on which we can form portfolios to hold for the next month. For each firm i , I compute a disaster- β from daily data using the following regression:

$$r_{it} = a_i + \sum_{k=0}^2 \beta_{ik} \Delta p_{RD,t-k}^* + e_{it}$$

The regression uses daily data from the previous quarter (i.e. t to $t - 90$). I define the disaster- β for each firm and at the end of each month as the sum of the estimated regression coefficients, i.e. $\beta_i = \sum_{k=0}^2 \beta_{ik}$. I compute disaster- β 's using daily data only for firms with a full set of daily returns over the previous quarter.

B.9 Computing Weekly and Monthly Returns Using Daily CRSP Data

B.10 Weekly Returns

In order to compute disaster- β using weekly data, I need weekly returns for each firm. To do so, I simply accumulate the daily returns within each week to compute a weekly return. To be considered in this accumulation, a firm must have at least 3 observations within each week.

B.11 Monthly Returns

To track portfolio performance, I need a monthly return of each firm. CRSP provides a monthly return, but return computations are calculated using end of month prices (i.e. the February return is the return from the end of the previous January to the end of February). This computation would be fine if portfolios were formed at the end of each month. However, to really ensure tradability, I estimate disaster- β 's at the end of each month, but then I actually form the portfolios using these betas on the next business day.

For example, suppose we are standing in January 1997. On January 31, 1997, I use the previous year's worth of data to estimate disaster- β 's. Then, on February 3, 1997 — the first business day after January 31 — I form the portfolios. I hold these until February 28, 1997, and then repeat the procedure.

I therefore use CRSP Daily returns to compute a within-month return measure for each firm in CRSP. To do so, I simply compound daily returns within each month. To be included in the analysis, the last observation for a firm in a month must be within 5 days of the end of the business month. In addition, each firm's first observation in a month must be within 5 days of the first of the business month.

C Appendix: Model Solution and Proofs

When presenting the model in Appendix A, I use the term “approximate” in describing expressions for 1) the stochastic discount factor, 2) price dividend ratios for the claim to aggregate consumption, and 3) price dividend ratios for claims to an arbitrary stock. In this appendix I will be precise in describing what is meant by approximate, and also solve the model exactly. In addition, I provide proofs of the theorems and propositions in the main text and in Appendix A.

C.1 The Claim to Aggregate Consumption

As is by now standard, a claim to aggregate consumption pays consumption as a dividend. Denote its price at time t as P_{Ct} . First, I define the resilience of this claim, H_{Ct} , as:

$$H_{Ct} := p_t \mathbb{E}_{t^-} \left[B_t^{1-\gamma} - 1 \right] - \gamma(1-\gamma) \frac{\sigma_C^2}{2} \quad (22)$$

As with an arbitrary stock, I model the resilience of the consumption claim as $H_{Ct} = H_{C^*} + \hat{H}_{Ct}$, where H_{C^*} is the long run mean and \hat{H}_{Ct} is deviations from that mean. Specifically,

$$\hat{H}_{Ct} = k_t + \frac{1-\chi}{2} v^2(k_t)$$

with the process k_t evolving as follows:

$$dk_t = -(\phi_{H,C} + \chi k_t) k_t dt + (\delta_C + \phi_{H,C} + \chi k_t) [(\chi - 1) v^2(k_t) dt + v(k_t) dW_{k,t}] \quad (23)$$

where $W_{k,t}$ is a standard Brownian motion. $v(k_t)$ is an arbitrary function that dictates the volatility of k_t and I require only that it goes to zero just before k_t hits $-\phi_C/\chi$. Note that this restriction is not much different than those put a standard [Cox, Ingersoll, and Ross \(1985\)](#) process. k_t is instead an LG-process that mean reverts at a speed $\phi_{H,C}$. The extra dt terms are small in magnitude and are a modeling convenience that comes with the LG class.

A Comment on the Modeling Approach and How I Use “Approximate” in the Main Text At this juncture, it may appear confusing why I don’t simply model H_{Ct} as an LG process directly, instead of modeling it as a function of the LG-process k_t . The reason is that I eventually want to have a price-dividend ratio for the consumption claim that is linear in an LG process. It turns out that in order for all expressions and processes to be internally consistent, there is little freedom from a modeling perspective. Following with [Campbell and Cochrane \(1999\)](#) and [Gabaix \(2012\)](#), I reverse engineered the process for k_t and the definition of H_{Ct} to ensure internal consistency. In addition, [Gabaix \(2012\)](#) shows that the volatility terms are small compared to the “leading” term in k_t , meaning $\hat{H}_{Ct} \simeq k_t$, and I adopt this argument here as well. This is the sense in which I use the phrase “approximate” in the main text. Notice also that in the power utility case, $\chi = 1$, and all of the expressions in the main text become exact.

Guess and Verify Approach I take a “guess and verify” approach to calculate the price-dividend ratio of the unlevered consumption claim. That is, I will first conjecture an equilibrium price dividend ratio for the aggregate consumption claim. From this conjecture, I will use the stochastic discount factor to derive the value of the aggregate consumption claim, and thus confirm that my original conjecture is true. Along the way, the following Lemma will prove useful:

Lemma 1. *The continuous time Epstein-Duffie-Zin stochastic discount factor evolves according to:*

$$\frac{dM_t}{M_{t^-}} = -\frac{\rho}{\chi} dt - \frac{1}{\chi\psi} \cdot \frac{dC_t}{C_{t^-}} + [(1 + dR_{C,t})^{1/\chi-1} - 1]$$

Proof. The proof follows from taking the limit of Equation (5) as $\tau \searrow 0$.³⁷ □

Next, I conjecture a linear form for the price-dividend ratio of the consumption claim.

Conjecture 1. *The price-dividend ratio of the unlevered consumption claim is given by:*

$$\frac{P_{C,t}}{C_t} = \frac{1}{\delta_C} \left(1 + \frac{\chi k_t}{\phi_{H,C} + \delta_C} \right), \quad \delta_C := \delta - g_C - \chi H_{C^*} \quad (24)$$

where $\delta := \rho + g_C/\psi$.

Let us now check the internal consistency of this conjecture. The instantaneous return to the aggregate consumption claim is:

$$\begin{aligned} dR_{C,t} &= \frac{dP_{C,t}}{P_{C,t}} + \frac{C_t}{P_{C,t}} dt \\ &= \frac{dC_t}{C_{t^-}} + \frac{d(P_{C,t}/C_t)}{P_{C,t}/C_t} + \frac{C_t}{P_{C,t}} dt \\ &= \frac{dC_t}{C_{t^-}} + \frac{\frac{\chi}{\phi_{H,C} + \delta_C} dk_t}{1 + \frac{\chi k_t}{\phi_{H,C} + \delta_C}} + \left(\frac{\delta_C}{1 + \frac{\chi k_t}{\phi_{H,C} + \delta_C}} - \delta_C \right) dt + \delta_C dt \\ &= \frac{dC_t}{C_{t^-}} + \delta_C dt + \frac{\chi dk_t}{\delta_C + \phi_{H,C} + \chi k_t} - \frac{\delta_C \chi k_t dt}{\delta_C + \phi_{H,C} + \chi k_t} \\ &= \frac{dC_t}{C_{t^-}} + [\delta_C - \chi k_t + \chi(\chi - 1)v^2(k_t)] dt + \chi v(k_t) dW_{k,t} \end{aligned} \quad (25)$$

where the last line comes from substituting in for dk_t and canceling like terms. For now, define $\mu_{kt} = [\delta_C - \chi k_t + \chi(\chi - 1)v^2(k_t)]$. Next, substitute Equation (25) and apply Lemma 1 to compute the evolution

³⁷For instance, using a Taylor expansion around 1:

$$\left(\frac{C_{t+\tau}}{C_t} \right)^{-1/\chi\psi} \approx 1 - \frac{1}{\chi\psi} \left(\frac{C_{t+\tau}}{C_t} - 1 \right) + \frac{1}{2} \times \frac{1}{\chi\psi} \left(\frac{1}{\chi\psi} + 1 \right) \left(\frac{C_{t+\tau}}{C_t} - 1 \right)^2$$

Then when taking continuous time limits, $C_{t+\tau}/C_t - 1$ converges to dC_t/C_t . The second order term is necessary to account for the standard Ito correction term that appears in stochastic calculus.

of the stochastic discount factor outside of disasters (“no disasters”):

$$\frac{dM_t}{M_t}|_{ND+1} = \left(1 - \frac{\rho}{\chi} dt\right) \times \left(1 - \frac{1}{\chi\psi} \frac{dC_t}{C_t} + \frac{1}{2} \times \frac{1}{\chi\psi} \left(\frac{1}{\chi\psi} + 1\right) \sigma_C^2 dt\right) \times (1 + dR_{C,t})^{1/\chi-1} \quad (26)$$

Our next task is to insert the expression derived for $dR_{C,t}$ in Equation (25) into the stochastic discount factor. To do so, we can Taylor expand the last line around zero, and as is standard in stochastic calculus, eliminate terms of $\mathcal{O}(dt^2)$ and higher:

$$\begin{aligned} \frac{dM_t}{M_t}|_{ND+1} &= \left(1 - \frac{\rho}{\chi} dt\right) \times \left(1 - \frac{1}{\chi\psi} \frac{dC_t}{C_t}|_{ND} + \frac{1}{2} \times \frac{1}{\chi\psi} \left(\frac{1}{\chi\psi} + 1\right) \sigma_C^2 dt\right) \times (1 + dR_{C,t})^{1/\chi-1} \\ &= \left(1 - \frac{\rho}{\chi} dt\right) \times \left(1 - \frac{1}{\chi\psi} \frac{dC_t}{C_t}|_{ND} + \frac{1}{2} \times \frac{1}{\chi\psi} \left(\frac{1}{\chi\psi} + 1\right) \sigma_C^2 dt\right) \times \\ &\quad \left[1 + \left(\frac{1}{\chi} - 1\right) \left(\frac{dC_t}{C_t}|_{ND} + \mu_{k_t} dt + \chi v(k_t) dW_{k,t}\right) + \frac{1}{2} \left(\frac{1}{\chi} - 1\right) \left(\frac{1}{\chi} - 2\right) (\sigma_C^2 + \chi^2 v^2(k_t)) dt\right] \\ &= \left(1 - \frac{\rho}{\chi} dt - \frac{1}{\chi\psi} \frac{dC_t}{C_t}|_{ND} + \frac{1}{2} \times \frac{1}{\chi\psi} \left(\frac{1}{\chi\psi} + 1\right) \sigma_C^2 dt\right) \times \\ &\quad \left[1 + \left(\frac{1}{\chi} - 1\right) \left(\frac{dC_t}{C_t}|_{ND} + \mu_{k_t} dt + \chi v(k_t) dW_{k,t}\right) + \frac{1}{2} \left(\frac{1}{\chi} - 1\right) \left(\frac{1}{\chi} - 2\right) (\sigma_C^2 + \chi^2 v^2(k_t)) dt\right] \\ &= 1 + \left(\frac{1}{\chi} - 1\right) \left(\frac{dC_t}{C_t}|_{ND} + \mu_{k_t} dt + \chi v(k_t) dW_{k,t}\right) + \frac{1}{2} \left(\frac{1}{\chi} - 1\right) \left(\frac{1}{\chi} - 2\right) (\sigma_C^2 + \chi^2 v^2(k_t)) dt + \\ &\quad - \frac{\rho}{\chi} dt - \frac{1}{\chi\psi} \frac{dC_t}{C_t}|_{ND} - \frac{1}{\chi\psi} \left(\frac{1}{\chi} - 1\right) \sigma_C^2 dt + \frac{1}{2} \times \frac{1}{\chi\psi} \left(\frac{1}{\chi\psi} + 1\right) \sigma_C^2 dt \\ \frac{dM_t}{M_t}|_{ND} &= -\frac{\rho}{\chi} dt + \left(\frac{1}{\chi} - 1\right) \mu_{k_t} dt + \left(\frac{1}{\chi} - 1 - \frac{1}{\chi\psi}\right) \frac{dC_t}{C_t}|_{ND} + (1 - \chi) v(k_t) dW_{k,t} \\ &\quad + \frac{1}{2} \left(\frac{1}{\chi} - 1\right) \left(\frac{1}{\chi} - 2\right) \chi^2 v^2(k_t) dt + \\ &\quad + \frac{\sigma_C^2}{2} \left[\left(\frac{1}{\chi} - 1\right) \left(\frac{1}{\chi} - 2\right) - \frac{2}{\chi\psi} \left(\frac{1}{\chi} - 1\right) + \frac{1}{\chi\psi} \left(\frac{1}{\chi\psi} + 1\right)\right] dt \end{aligned} \quad (27)$$

where I eliminated terms containing $dW_{C,t} dW_{k,t}$ because I assume the two shocks are independent. Tedious algebra shows that the term multiplying σ_C^2 can be dramatically simplified:

$$\left[\left(\frac{1}{\chi} - 1\right) \left(\frac{1}{\chi} - 2\right) - \frac{2}{\chi\psi} \left(\frac{1}{\chi} - 1\right) + \frac{1}{\chi\psi} \left(\frac{1}{\chi\psi} + 1\right)\right] = \gamma(1 + \gamma)$$

Finally, substituting in for $\mu_{k_t} = [\delta_C - \chi k_t + \chi(\chi - 1)v^2(k_t)]$ and k_t yields, after some simplification:

$$\begin{aligned} \frac{dM_t}{M_t} \Big|_{ND} &= (-\delta + (\chi - 1)H_C^*)dt \\ &\quad + \frac{\sigma_C^2}{2}\gamma(1 + \gamma)dt - \gamma\sigma_C dW_{C,t} \\ &\quad + (\chi - 1) \left[\left(k_t + \frac{v^2(k_t)}{2} \right) dt - v(k_t)dW_{k,t} \right] \end{aligned} \quad (28)$$

where I have used the identities $\delta = \rho - g_c/\psi$ and $\delta_c = \delta - g_c - \chi H_C^*$.

Note that in the power utility model, $\chi = 1$, which means the stochastic discount factor outside of disasters takes its usual form (e.g. Cochrane (2005)). Further notice if there is a disaster that M_t is simply multiplied by $B_t^{-1/(\chi\psi)+1/\chi-1} = B_t^{-\gamma}$, which means we can write the SDF as:

$$\frac{dM_t}{M_t} = \frac{dM_t}{M_t} \Big|_{ND} + (B_t^{-\gamma} - 1) dJ_t \quad (29)$$

We are now in a position to verify Conjecture 1 using the LG machinery. From Equation (29) and the law of motion for consumption, I calculate:

$$\begin{aligned} \frac{d(M_t C_t)}{M_t C_t} &= (-\delta + (\chi - 1)H_C^*)dt - \gamma\sigma_C dW_{C,t} + (\chi - 1) \left[\left(k_t + \frac{v^2(k_t)}{2} \right) dt - v(k_t)dW_{k,t} \right] \\ &\quad + \frac{\sigma_C^2}{2}\gamma(1 + \gamma)dt + g_C dt + \sigma_C dW_{C,t} - \gamma\sigma_C^2 dt + (B_t^{1-\gamma} - 1) dJ_t \end{aligned}$$

Thus, the first LG moment is:

$$\begin{aligned} \mathbb{E}_t \left[\frac{d(M_t C_t)}{M_t C_t} \right] / dt &= -\delta + (\chi - 1) \left(H_C^* + k_t + \frac{v^2(k_t)}{2} \right) + g_C + \frac{\sigma_C^2}{2}\gamma(1 + \gamma) - \gamma\sigma_C^2 + p_t \mathbb{E}_t [B_t^{1-\gamma} - 1] \\ &= -\delta + (\chi - 1) \left(H_C^* + k_t + \frac{v^2(k_t)}{2} \right) + g_C + H_{C_t} \\ &= -\delta + (\chi - 1) \left(H_C^* + k_t + \frac{v^2(k_t)}{2} \right) + g_C + H_C^* + \hat{H}_{C_t} \\ &= -\delta + (\chi - 1) \left(H_C^* + k_t + \frac{v^2(k_t)}{2} \right) + g_C + H_C^* + k_t + \frac{1-\chi}{2} v^2(k_t) \\ &= -\delta + \chi H_C^* + g_C + \chi k_t \\ &= -\delta_C + \chi k_t \end{aligned} \quad (30)$$

Let's compute the second LG moment:

$$\begin{aligned}
\mathbb{E}_t \left[\frac{d(M_t C_t k_t)}{M_t C_t} \right] / dt &= \mathbb{E}_t \left[\frac{d(M_t C_t)}{M_t C_t} \right] / dt \cdot k_t + \mathbb{E}_t [dk_t] / dt + \left\langle \frac{dM_t}{M_t}, dk_t \right\rangle \\
&= (-\delta_C + \chi k_t) k_t - (\phi_{H,C} + \chi k_t) k_t dt + (\delta_C + \phi_{H,C} + \chi k_t)(\chi - 1) v^2(k_t) dt \\
&\quad + (\delta_C + \phi_{H,C} + \chi k_t)(1 - \chi) v^2(k_t) dt \\
&= -(\delta_C + \phi_{H,C}) k_t
\end{aligned}$$

Finally, applying Theorem 4 of [Gabaix \(2009\)](#) yields the price-dividend of the consumption claim:

$$\frac{P_{Ct}}{C_t} = \frac{1}{\delta_C} \left(1 + \frac{\chi k_t}{\phi_{H,C} + \delta_C} \right)$$

So, indeed Conjecture 1 is correct. It is also worth noting that when the IES is one (i.e. $\psi = 1$), then the price-dividend ratio of the consumption claim collapses to a constant of ρ^{-1} as in Weil (1990). Notice also that if we shut down time-varying disaster risks, then the dividend-yield of the consumption claim becomes:

$$\begin{aligned}
\frac{C_t}{P_{Ct}} &= \delta_C \\
&= \rho + \left(\frac{1}{\psi} - 1 \right) g_C - \frac{1}{2} \gamma \left(\frac{1}{\psi} - 1 \right) \sigma_c^2 - \frac{1/\psi - 1}{\gamma - 1} p(\mathbb{E}[B^{1-\gamma}] - 1)
\end{aligned}$$

This expression matches the derivation of Equation (13) in [Barro \(2009\)](#).

C.2 Equilibrium Valuation of Any Arbitrary Stock

Similar to the aggregate consumption claim, define the resilience of firm i 's stock as:

$$H_{it} := p_t \mathbb{E}_t [B_t^{-\gamma} F_{i,t} - 1] + (\chi - 1) \left(H_{Ct} + \frac{\chi v^2(k_t)}{2} \right) + \frac{\sigma_C^2}{2} \gamma (1 + \gamma) \quad (31)$$

Further suppose that $H_{it} = H_{i*} + \hat{H}_{it}$, and the variable part of firm-level resilience follows an LG process (repeated from the main text):

$$d\hat{H}_{it} + \left\langle \frac{dM_t}{M_t} + \frac{dD_{it}}{D_{it}}, dN_{it}^H \right\rangle^{ND} dt = -(\phi_{H,i} + \hat{H}_{it}) \hat{H}_{it} dt + dN_{it}^H \quad (32)$$

where again N_{it}^H is a mean-zero shock to firm-level resilience and $\langle x, y \rangle^{ND}$ denotes covariance of x and y , conditional on no disasters.

Next, to compute the price-dividend ratio, we need to verify the conditions for H_{it} to be used in

conjunction with the LG machinery. Start with the first LG moment:

$$\begin{aligned} \frac{d(M_t D_{it})}{M_t^- D_{it}^-} &= (-\delta + (\chi - 1)H_C^*)dt - \gamma \sigma_C dW_{C,t} + (\chi - 1) \left[\left(k_t + \frac{v^2(k_t)}{2} \right) dt - v(k_t) dW_{k,t} \right] \\ &\quad + \frac{\sigma_C^2}{2} \gamma (1 + \gamma) dt + g_{iD} dt + dN_{it}^D + \left\langle \frac{dM_t}{M_t}, \frac{dD_{it}}{D_{it}} \right\rangle^{ND} \\ &\quad + (B_{t+dt}^{-\gamma} F_{i,t+dt} - 1) dJ_t \end{aligned}$$

Using the assumption that shocks to dividend growth in normal times, dN_{it}^D , are independent of all other shocks gives us:

$$\begin{aligned} \mathbb{E}_t \left[\frac{d(M_t D_{it})}{M_t^- D_{it}^-} \right] / dt &= -\delta + (\chi - 1) \left(H_C^* + k_t + \frac{v^2(k_t)}{2} \right) + g_{iD} + \frac{\sigma_C^2}{2} \gamma (1 + \gamma) dt + p_t \mathbb{E}_t (B_t^{-\gamma} F_{i,t} - 1) \\ &= -\delta + (\chi - 1) \left(H_C^* + k_t + \frac{1 - \chi}{2} v^2(k_t) - \frac{1 - \chi}{2} v^2(k_t) + \frac{v^2(k_t)}{2} \right) \\ &\quad + g_{iD} + \frac{\sigma_C^2}{2} \gamma (1 + \gamma) dt + p_t \mathbb{E}_t (B_{t+dt}^{-\gamma} F_{i,t+dt} - 1) \\ &= -\delta + g_{iD} + (\chi - 1) \left(H_C^* + \frac{\chi v^2(k_t)}{2} \right) + \frac{\sigma_C^2}{2} \gamma (1 + \gamma) dt + p_t \mathbb{E}_t (B_t^{-\gamma} F_{i,t} - 1) \\ &= -\delta + g_{iD} + H_{it} \\ &= -\delta + g_{iD} + H_i^* + \hat{H}_{it} \\ &= -\delta_i + \hat{H}_{it} \end{aligned}$$

where I define $\delta_i := \delta - g_{iD} - H_i^*$. Now, the second LG moment:

$$\begin{aligned} \mathbb{E}_t \left[\frac{d(M_t D_{it} \hat{H}_{it})}{M_t^- D_{it}^-} \right] / dt &= (-\delta_i + \hat{H}_{it}) \hat{H}_{it} + \left\langle \frac{dM_t}{M_t} + \frac{dD_{it}}{D_{it}}, d\hat{H}_{it} \right\rangle + \mathbb{E}_t [d\hat{H}_{it}] / dt \\ &= (-\delta_i + \hat{H}_{it}) \hat{H}_{it} - (\phi_H + \hat{H}_{it}) \hat{H}_{it} \\ &= -(\delta_i + \phi_H) \hat{H}_{it} \end{aligned}$$

Finally, using Theorem 4 of Gabaix (2012), we get the equilibrium price-dividend ratio:

$$\frac{P_{it}}{D_{it}} = \frac{1}{\delta_i} \left(1 + \frac{\hat{H}_{it}}{\delta_i + \phi_{H,i}} \right) \quad (33)$$

C.3 Equilibrium Returns

From Equation (33), it is straightforward to derive the equilibrium return process for the firm. Start with the instantaneous change in the level of prices:

$$dP_{it} = \frac{1}{\delta_i} \left[dD_{it} + \frac{1}{\delta_i + \phi_{H,i}} (dD_{it} \hat{H}_{it} + D_{it} d\hat{H}_{it} + \langle dD_{it}, d\hat{H}_{it} \rangle) \right]$$

so ex-dividend returns are,

$$\begin{aligned}
\frac{dP_{it}}{P_{it}} &= \frac{\frac{1}{\delta_i} \left[dD_{it} + \frac{1}{\delta_i + \phi_{H,i}} (dD_{it} \hat{H}_{it} + D_{it} d\hat{H}_{it} + \langle dD_{it}, d\hat{H}_{it} \rangle) \right]}{\frac{D_{it}}{\delta_i} \left(1 + \frac{\hat{H}_{it}}{\delta_i + \phi_{H,i}} \right)} \\
&= \frac{dD_{it} + \frac{1}{\delta_i + \phi_{H,i}} (dD_{it} \hat{H}_{it} + D_{it} d\hat{H}_{it})}{D_{it} \left(1 + \frac{\hat{H}_{it}}{\delta_i + \phi_{H,i}} \right)} \\
&= \left(1 + \frac{\hat{H}_{it}}{\delta_i + \phi_{H,i}} \right)^{-1} \left[\frac{dD_{it}}{D_{it}} + \frac{dD_{it}}{D_{it}} \left(\frac{\hat{H}_{it}}{\delta_i + \phi_{H,i}} \right) + \frac{d\hat{H}_{it}}{\delta_i + \phi_{H,i}} \right] \\
&= \left(\frac{\delta_i + \phi_{H,i}}{\delta_i + \phi_{H,i} + \hat{H}_{it}} \right) \left[\frac{dD_{it}}{D_{it}} \left(1 + \frac{\hat{H}_{it}}{\delta_i + \phi_{H,i}} \right) + d\hat{H}_{it} \left(\frac{1}{\delta_i + \phi_{H,i}} \right) \right] \\
&= \frac{dD_{it}}{D_{it}} \left(\frac{\delta_i + \phi_{H,i}}{\delta_i + \phi_{H,i} + \hat{H}_{it}} \right) \left(\frac{\delta_i + \phi_{H,i} + \hat{H}_{it}}{\delta_i + \phi_{H,i}} \right) + \frac{d\hat{H}_{it}}{\delta_i + \phi_{H,i} + \hat{H}_{it}} \\
&= \frac{dD_{it}}{D_{it}} + \frac{d\hat{H}_{it}}{\delta_i + \phi_{H,i} + \hat{H}_{it}} \tag{34}
\end{aligned}$$

Substituting Equation (6) and (32) into (34) yields:

$$\begin{aligned}
\frac{dP_{it}}{P_{it}} &= (g_{iD} dt + dN_{it}^D + (F_{it} - 1) d\mathcal{J}_t) + \\
&\quad \left(\frac{1}{\delta_i + \phi_{H,i} + \hat{H}_{it}} \right) \left(-(\phi_{H,i} + \hat{H}_{it}) \hat{H}_{it} dt + dN_{it}^H - \left\langle \frac{dM_t}{M_t} + \frac{dD_{it}}{D_{it}}, dN_{it}^H \right\rangle^{ND} dt \right) \\
&= \mu_{it} dt + \sigma_{it} dW_{it} + (F_{it} - 1) dJ_t \tag{35}
\end{aligned}$$

where

$$\begin{aligned}
\mu_{it} &= g_{iD} - \left(\frac{(\phi_{H,i} + \hat{H}_{it}) \hat{H}_{it} + \left\langle \frac{dM_t}{M_t} + \frac{dD_{it}}{D_{it}}, dN_{it}^H \right\rangle^{ND}}{\delta_i + \phi_{H,i} + \hat{H}_{it}} \right) \\
\sigma_{it} &= \sqrt{\text{var}_t(dN_{it}^D) + \left(\frac{1}{\delta_i + \phi_{H,i} + \hat{H}_{it}} \right)^2 \text{var}_t(dN_{it}^H)}
\end{aligned}$$

where dW_{it} is a standard Wiener process.

Where did the dW_{it} in Equation (35) come from? Recall that I only restricted N_{it}^D and N_{it}^H to be mean-zero innovations. Importantly, this was all that was necessary to derive equilibrium price-dividend ratios as in Equation (33). The remaining task is to engineer processes for N_t^D and N_t^H that are based on Wiener processes, have mean zero increments, and also satisfy the restrictions on variance that are required for LG processes. Here is a sketch of one way to accomplish this: set $dN_t^D = \sigma_{it}^D dW_{it}^D$ for an arbitrary volatility process, σ_{it}^D . Then, define $dN_{it}^H = -a N_{it}^H - \alpha (N_{it}^H - \lambda) dW_{it}^H$. This is the process used by [Menzly, Santos, and Veronesi \(2004\)](#). The parameter a determines the mean-reversion speed of the process and λ dic-

tates the domain of the process so that $N_{it}^H \in [\lambda, \infty)$. Finally, set λ equal to the appropriate lower bound as dictated by the LG restrictions.³⁸

It is also important to recognize that dW_{it} is not necessarily independent of other shocks in the model. Namely, given the definition of firm resilience, dW_{it} is mechanically correlated with dW_{kt} . As such, we can decompose dW_{it} as follows: $dW_{it} = \rho dW_{kt} + \sqrt{1 - \rho^2} d\widetilde{W}_{it}$, where $d\widetilde{W}_{it}$ is an independent Brownian motion, and ρ is a correlation coefficient that is implicitly determined by the definition of resilience (both the consumption claim and at the firm level). Putting this all together, we obtain the following internally consistent process for firm-level returns:

$$\frac{dP_{it}}{P_{it}} = \mu_{it} dt + \sigma_{it} \left(\rho dW_{kt} + \sqrt{1 - \rho^2} d\widetilde{W}_{it} \right) + (F_{it} - 1) dJ_t \quad (36)$$

C.4 Risk-Neutral Price Dynamics

I solve for the risk-neutral price dynamics of any stock in three steps: (i) solve for the equilibrium riskless short rate process, r_t ; (ii) define the Radon Nikodym derivative from the stochastic discount factor process; (iii) use Girsanov's theorem for jumps to solve for risk-neutral price dynamics.

Step (i) is necessary in order to pin down the risk-neutral drift (and thus the excess risk premium) as a function of the model's primitives. One reference for Girsanov's theorem for jumps can be found in Björk, Kabanov, and Runggaldier (1997).

³⁸Footnote 9 in Gabaix [Gabaix \(2012\)](#) details these restrictions based on how resilience is defined in the disaster model.

C.4.1 Equilibrium Short Rate Process

Let me first hone down the (dynamic) short rate r_t ,

$$\begin{aligned}
r_{ft} &= \mathbb{E}_t \left[-\frac{dM_t}{M_t} \right] / dt \\
&= (\delta - (\chi - 1)H_{C^*}) - (\chi - 1) \left[k_t + \frac{v^2(k_t)}{2} \right] - \frac{\sigma_C^2}{2} \gamma(1 + \gamma) - p_t \mathbb{E}_t [B_t^{-\gamma} - 1] \\
&= \delta - (\chi - 1) \left[H_{C^*} + k_t + \frac{v^2(k_t)(1 - \chi)}{2} + \frac{\chi v^2(k_t)}{2} \right] - \frac{\sigma_C^2}{2} \gamma(1 + \gamma) - p_t \mathbb{E}_t [B_t^{-\gamma} - 1] \\
&= \delta - (\chi - 1) \left[H_{C_t} + \frac{\chi v^2(k_t)}{2} \right] - \frac{\sigma_C^2}{2} \gamma(1 + \gamma) - p_t \mathbb{E}_t [B_t^{-\gamma} - 1] \\
&= \delta - (\chi - 1) \left[p_t \mathbb{E}_t [B_t^{1-\gamma} - 1] - \gamma(1 - \gamma) \frac{\sigma_C^2}{2} + \frac{\chi v^2(k_t)}{2} \right] - \frac{\sigma_C^2}{2} \gamma(1 + \gamma) - p_t \mathbb{E}_t [B_t^{-\gamma} - 1] \\
&\quad \delta - \frac{\sigma_C^2}{2} \gamma [\chi(\gamma - 1) + 2] - p_t \mathbb{E}_t [B_t^{-\gamma} (1 + (\chi - 1)B_t) - \chi] - \frac{(\chi - 1)\chi v^2(k_t)}{2} \\
&= \delta - \frac{\sigma_C^2}{2} \gamma \left[\frac{1}{\psi} + 1 \right] - p_t \mathbb{E}_t [B_t^{-\gamma} (1 + (\chi - 1)B_t) - \chi] - \frac{(\chi - 1)\chi v^2(k_t)}{2} \tag{37}
\end{aligned}$$

$$\approx \delta - \frac{\sigma_C^2}{2} \gamma \left[\frac{1}{\psi} + 1 \right] - p_t \mathbb{E}_t [B_t^{-\gamma} (1 + (\chi - 1)B_t) - \chi] \tag{38}$$

where the approximation in the last line comes from ignoring the higher order $v^2(k_t)$ term.

Notice that if $\psi = 1/\gamma$ as in the power utility case, then we get the exact expression in [Barro \(2009\)](#) for time-invariant disasters and [Gabaix \(2012\)](#) for time-varying disasters. Furthermore, if $\psi = 1$, then the expression matches the risk-free rate derived in [Wachter \(2013\)](#).

C.4.2 Defining the Radon-Nikodym Derivative

Define the process Z_t as follows:

$$Z_t := \exp \left(\int_0^t r_{fs} ds \right) M_t \tag{39}$$

so that,

$$\begin{aligned}
dZ_t &= \exp\left(\int_0^t r_{f_s} ds\right) [r_{f_t} dt] M_t + \exp\left(\int_0^t r_{f_s} ds\right) dM_t \\
\frac{dZ_t}{Z_t} &= r_{f_t} dt + \frac{dM_t}{M_t} \\
&= (\delta - (\chi - 1)H_C^*) dt - (\chi - 1) \left[k_t + \frac{v^2(k_t)}{2} \right] dt - \frac{\sigma_C^2}{2} \gamma(1 + \gamma) dt - p_t \mathbb{E}_t [B_t^{-\gamma} - 1] dt \\
&\quad + (-\delta + (\chi - 1)H_C^*) dt + \frac{\sigma_C^2}{2} \gamma(1 + \gamma) dt - \gamma \sigma_C dW_{C,t} \\
&\quad + (\chi - 1) \left[\left(k_t + \frac{v^2(k_t)}{2} \right) dt - v(k_t) dW_{k,t} \right] + (B_t^{-\gamma} - 1) dJ_t \\
&= -\gamma \sigma_C dW_{C,t} + (B_t^{-\gamma} - 1) dJ_t - p_t \mathbb{E}_t [B_t^{-\gamma} - 1] dt - (\chi - 1) v(k_t) dW_{k,t}
\end{aligned} \tag{40}$$

It should be clear from Equation (40) that Z_t is a martingale under the physical measure, \mathbb{P} .

Next, define risk-neutral measure, \mathbb{Q} , via its Radon-Nikodym derivative:

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = \frac{Z_t}{Z_0} \tag{41}$$

C.4.3 Girsanov's Theorem to Change Measure

Girsanov's theorem for jumps, e.g. Kabanov, and Runggaldier (1997), states that the $dW_{C_t}^* \equiv dW_{C_t} + \gamma \sigma_C dt$ and $dW_{k_t}^* = dW_{k_t} + (\chi - 1)v(k_t)dt$ are both Brownian motions under the new measure \mathbb{Q} . In addition, the disaster arrival rate p_t is shifted so that disasters dJ_t arrive at a risk-neutral rate $p_t^* \equiv p_t \mathbb{E}_t [B_t^{-\gamma}]$. Notice that when $B_t < 0$, $p_t^* > p_t$. Intuitively, the agent prices disasters as if they are more likely under the risk-neutral measure. The severity of disasters at the aggregate level also changes under the risk-neutral measure. Suppose that the probability distribution function of B_t is given by $f_{B_t}(x)$. Under the risk-neutral measure, Girsanov's theorem states that B_t is distributed according to $f_{B_t}^*(x)$, where $f_{B_t}^*(x)$ is defined as:

$$f_{B_t}^*(x) = \frac{x^{-\gamma} f_{B_t}(x)}{\mathbb{E}_t [B_t^{-\gamma}]}$$

It is easy to verify that $f_{B_t}^*(x)$ is a valid distribution function. Furthermore, $\mathbb{E}_t^{\mathbb{Q}} [B_t] = \mathbb{E}_t^{\mathbb{P}} [B_t^{1-\gamma}] / \mathbb{E}_t^{\mathbb{P}} [B_t^{-\gamma}]$. With uncertainty under the risk-neutral measure appropriately defined, we can now determine the price of a consumption claim.

Pricing the Consumption Claim:

From Equation (25), we know that the return to the consumption claim is:

$$\begin{aligned}
dR_{Ct} &= g_C dt + \sigma_C dW_{Ct} + [\delta_C - \chi k_t + \chi(\chi - 1)v^2(k_t)] dt + \chi v(k_t) dW_{kt} + (B_t - 1)dJ_t \\
&= \mu_{Ct} dt + \sigma_C dW_{Ct} + \chi v(k_t) dW_{kt} + (B_t - 1)dJ_t \\
&= \mu_{Ct} dt + \sigma_C [dW_{Ct}^* - \gamma \sigma_C dt] + \chi v(k_t) [dW_{kt}^* - (\chi - 1)v(k_t) dt] + (B_t - 1)dJ_t \\
&= \mu_{Ct} dt - \gamma \sigma_C^2 dt - \chi(\chi - 1)v^2(k_t) dt + \sigma_C dW_{Ct}^* + \chi v(k_t) dW_{kt}^* + (B_t - 1)dJ_t \\
&= \mu_{Ct} dt - \gamma \sigma_C^2 dt - \chi(\chi - 1)v^2(k_t) dt + \sigma_C dW_{Ct}^* + \chi v(k_t) dW_{kt}^* + (B_t - 1)dJ_t \\
&\quad + p_t^* \mathbb{E}_t^{\mathbb{Q}} [B_t - 1] - p_t^* \mathbb{E}_t^{\mathbb{Q}} [B_t - 1] \\
&= \mu_{Ct} dt - \gamma \sigma_C^2 dt - \chi(\chi - 1)v^2(k_t) dt + \sigma_C dW_{Ct}^* + \chi v(k_t) dW_{kt}^* + (B_t - 1)dJ_t \\
&\quad + p_t \mathbb{E}_t^{\mathbb{P}} [B_t^{1-\gamma} - B_t^{-\gamma}] dt - p_t^* \mathbb{E}_t^{\mathbb{Q}} [B_t - 1] dt
\end{aligned}$$

where $\mu_{Ct} = g_C + [\delta_C - \chi k_t + \chi(\chi - 1)v^2(k_t)]$. Under no arbitrage, we know that under the risk-neutral measure the expected change in the consumption claim must equal the risk-free rate (i.e. $\mathbb{E}_t^{\mathbb{Q}} [dR_{Ct}] = r_{ft} dt$):

$$\begin{aligned}
\mu_{Ct} dt - \gamma \sigma_C^2 dt - \chi(\chi - 1)v^2(k_t) dt + p_t \mathbb{E}_t^{\mathbb{P}} [B_t^{1-\gamma} - B_t^{-\gamma}] dt &= r_{ft} dt \\
&\Leftrightarrow \\
\mu_{Ct} - r_{ft} &= \gamma \sigma_C^2 + \chi(\chi - 1)v^2(k_t) + p_t \mathbb{E}_t^{\mathbb{P}} [B_t^{-\gamma}(1 - B_t)]
\end{aligned} \tag{42}$$

The expected return on the consumption claim is the sum of three components. The first, $\gamma \sigma_C^2$, is the standard result from power-utility economies, and we know that σ_C is too small to match expected unlevered stock market returns on its own. The second term $\chi(\chi - 1)v^2(k_t)$ arises from Epstein-Zin-Weil preferences. When the IES $\psi > 1$, then more uncertainty in resilience (the disaster likelihood of severity) leads to a higher equity premium. This result mirrors the findings of, for instance, Bansal and Yaron (2003). The last term reflects the additional premium required for bearing disaster risk.³⁹

Pricing an Arbitrary Stock:

³⁹Note that this is the expected return to the consumption claim, conditional on no disasters. This is because $\mathbb{E}_t^{\mathbb{P}} [dR_{Ct}] = \mu_{Ct} dt + p_t \mathbb{E}_t^{\mathbb{P}} [B_t - 1] dt$. The unconditional expected return of the consumption claim is thus given by: $\mathbb{E}_t^{\mathbb{P}} [dR_{Ct}] - r_{ft} dt = \gamma \sigma_C^2 + \chi(\chi - 1)v^2(k_t) + p_t \mathbb{E}_t^{\mathbb{P}} [(B_t^{-\gamma} - 1)(1 - B_t)] dt$.

Start from the return process for an individual firm:

$$\begin{aligned}
\frac{dP_{it}}{P_{it}} &= \mu_{it}dt + \sigma_{it} \left(\rho dW_{kt}^* - \rho(\chi - 1)v(k_t)dt + \sqrt{1 - \rho^2}d\widetilde{W}_{it} \right) + (F_{it} - 1)dJ_t \\
&= \mu_{it}dt + \sigma_{it}\sqrt{1 - \rho^2}d\widetilde{W}_{it} + \sigma_{it}\rho dW_{kt}^* - \rho(\chi - 1)v(k_t)\sigma_{it}dt \\
&\quad + p_t\mathbb{E}_t[B_t^{-\gamma}(F_{it} - 1)]dt - p_t\mathbb{E}_t[B_t^{-\gamma}(F_{it} - 1)]dt + (F_{it} - 1)dJ_t \\
&= [\mu_{it} + p_t\mathbb{E}_t[B_t^{-\gamma}(F_{it} - 1)] - \rho(\chi - 1)v(k_t)\sigma_{it}]dt \\
&\quad + \sigma_{it}\sqrt{1 - \rho^2}d\widetilde{W}_{it} + \sigma_{it}\rho dW_{kt}^* - p_t\mathbb{E}_t[B_t^{-\gamma}(F_{it} - 1)]dt + (F_{it} - 1)dJ_t \tag{43}
\end{aligned}$$

No-arbitrage implies that, under the \mathbb{Q} -measure, the discounted stock price is a Martingale. In order for this to be true, it must be the case that:

$$\begin{aligned}
\mu_{it} + p_t\mathbb{E}_t[B_t^{-\gamma}(F_{it} - 1)] - \rho(\chi - 1)v(k_t)\sigma_{it} &= r_{ft} \\
&\Leftrightarrow \\
\mu_{it} - r_{ft} &= \rho(\chi - 1)v(k_t)\sigma_{it} + p_t\mathbb{E}_t[B_t^{-\gamma}(1 - F_{it})] \\
&\approx p_t\mathbb{E}_t[B_t^{-\gamma}(1 - F_{it})] \tag{44}
\end{aligned}$$

where the approximation comes from the fact that $v(k_t)$ is in practice very small.⁴⁰ Equation (44) is the equilibrium risk premium for any stock.

Finally, we obtain the risk-neutral return process for the firm:

$$\begin{aligned}
\frac{dP_{it}^*}{P_{it}^*} &= r_{ft}dt - p_t\mathbb{E}_t[B_t^{-\gamma}(F_{it} - 1)]dt + \sigma_{it}\sqrt{1 - \rho^2}d\widetilde{W}_{it} + \sigma_{it}\rho dW_{kt}^* + (F_{it} - 1)dJ_t^* \\
&= r_{ft}dt - p_t\mathbb{E}_t[B_t^{-\gamma}(F_{it} - 1)]dt + \sigma_{it}dW_{it}^* + (F_{it} - 1)dJ_t^* \tag{45}
\end{aligned}$$

where $dW_{it}^* = \sqrt{1 - \rho^2}d\widetilde{W}_{it} + \rho dW_{kt}^*$. Here, I use the * superscript to denote the appropriate processes under \mathbb{Q} . The important takeaway of this exercise though is the risk-neutral jump process at the firm level.

C.5 Proof of Proposition 2

First notice that the risk-neutral probability of a disaster over the next Δ is $1 - e^{-pB^{-\gamma}\Delta} \approx pB^{-\gamma}\Delta = p^*$. Under the assumption that $\sigma_{it} = \sigma_i, p_t = p, B_t = B$, and $F_{it} = F_i$ are all constants, the risk-neutral stock price at time $t + \Delta$ is given by:

$$P_{i,t+\Delta} = \begin{cases} P_{it} \exp\{(\nu_i - \sigma_i^2/2)\Delta + \sigma_i u\sqrt{\Delta}\} & \text{if a disaster doesn't occur} \\ P_{it}F_i \exp\{(\nu_i - \sigma_i^2/2)\Delta + \sigma_i u\sqrt{\Delta}\} & \text{if a disaster occurs} \end{cases}$$

⁴⁰In addition, if independent innovations to firm-level recovery F_{it} are a large component of firm resilience, then ρ will also be very small.

where $v_i = r_f - pB^{-\gamma}(1 - F_i)$ and u is a standard normal random variable under the risk-neutral measure, $u \sim^{\mathbb{Q}} N(0, 1)$.

Proof. Next, recall that the Black-Scholes formula (with zero interest rate) for puts computes the following expectation:

$$\mathcal{P}_i^{BS}(S, K) = \mathbb{E}^M \left[\max \left(0, K - S e^{-\sigma_i^2/2 + \sigma_i z \sqrt{\Delta}} \right) \right]$$

where M is an arbitrary measure such that $z \sim^M N(0, 1)$ is a standard normal under that measure. For notational ease, I suppress the time to maturity Δ .

Now, in our context, the value of a put option is:

$$\begin{aligned} \mathcal{P}_i(P_{it}, K) &= e^{-r_f \Delta} \mathbb{E}^{\mathbb{Q}} [\max(0, K - P_{t+\Delta})] \\ &= e^{-r_f \Delta} \left\{ p^* \mathbb{E}_{No-Disas}^{\mathbb{Q}} [\max(0, K - P_{t+\Delta})] + (1 - p^*) \mathbb{E}_{Disas}^{\mathbb{Q}} [\max(0, K - P_{t+\Delta})] \right\} \end{aligned} \quad (46)$$

where $\mathbb{E}_{No-Disas}^{\mathbb{Q}}[\cdot]$ is the expectation under the risk-neutral measure, conditional on no disaster occurring. This is straightforward to compute:

$$\begin{aligned} \mathbb{E}_{No-Disas}^{\mathbb{Q}} [\max(0, K - P_{t+\Delta})] &= \mathbb{E}^{\mathbb{Q}} \left[\max(0, K - P_{it} \exp \{ (v_i - \sigma_i^2/2) \Delta + \sigma u \sqrt{\Delta} \}) \right] \\ &= e^{v_i \Delta} \mathbb{E}^{\mathbb{Q}} \left[\max(0, K - P_{it} \exp \{ (-\sigma_i^2/2) \Delta + \sigma u \sqrt{\Delta} \}) \right] \\ &= e^{v_i \Delta} \mathcal{P}_i^{BS}(P_{it}, K) \end{aligned}$$

It is straightforward to repeat the same argument to compute $\mathbb{E}_{Disas}^{\mathbb{Q}} [\max(0, K - P_{t+\Delta})]$. Thus, Equation (46) reduces to:

$$\begin{aligned} \mathcal{P}_i(P_{it}, K) &= e^{(v_i - r_f) \Delta} \left[p^* \times \mathcal{P}_i^{BS}(F_i P_{it}, K) + (1 - p^*) \times \mathcal{P}_i^{BS}(P_{it}, K) \right] \\ &= e^{(-p^*(1 - F_i)) \Delta} \left[p^* \times \mathcal{P}_i^{BS}(F_i P_{it}, K) + (1 - p^*) \times \mathcal{P}_i^{BS}(P_{it}, K) \right] \end{aligned}$$

which yields the desired result. A parallel argument holds for the call option value □

C.6 Proof of Theorem 3

Define $g(v) = \ln(v)$ and $\psi(x) = 1 + x + x^2/2 - e^x$. As defined in the main text, $f_i(F; t)$ is the density of F at time t . Finally, define $\Psi(f_i(F; t))$ as:

$$\Psi(f_i(F; t)) = 2 \int_{\mathbb{R}^+} \psi(g(v)) f(v; t) dv \quad (47)$$

The proof follows from applications of Proposition 1 of Carr and Wu (2009), and Proposition 2 of Du and Kapadia (2013).

In the case where F_i is constant, then Ψ takes a very simple form:

$$\Psi^c(F_i; \gamma) = 2 \left[1 + \ln(F_{it}) + [\ln(F_{it})]^2 / 2 - F_{it} \right] \quad (48)$$

C.7 Proof of Proposition 1

From Proposition 2, we know that the price dividend ratio of firm i at an arbitrary time $t + h$ is:

$$\frac{P_{i,t+h}}{D_{i,t+h}} = \frac{1}{\delta_i} \left(1 + \frac{\hat{H}_{i,t+h}}{\delta_i + \phi_{H,i}} \right)$$

Computing log-returns between $t + h$ and t is thus straightforward:

$$\begin{aligned} \log\left(\frac{P_{i,t+h}}{P_{i,t}}\right) &= \log\left(\frac{D_{i,t+h}}{D_{i,t}}\right) + \log\left(1 + \frac{\hat{H}_{i,t+h}}{\delta_i + \phi_{H,i}}\right) - \log\left(1 + \frac{\hat{H}_{i,t}}{\delta_i + \phi_{H,i}}\right) \\ &\approx \log\left(\frac{D_{i,t+h}}{D_{i,t}}\right) + \frac{1}{\delta_i + \phi_{H,i}} \times (\hat{H}_{i,t+h} - \hat{H}_{i,t}) \end{aligned} \quad (49)$$

where the approximation in the second line comes from a first-order Taylor expansion around $\hat{H} = 0$. This is a good approximation in this case because the mean of \hat{H} is by construction zero.

Next, using the definition of resilience $\hat{H}_{i,t}$, recognize that:

$$\begin{aligned} \hat{H}_{i,t+h} - \hat{H}_{i,t} &= p_{t+h} \times \{\mathbb{E}_{t+h,-} [B_{t+h}^{-\gamma} (F_{i,t+h} + (\chi - 1)B_{t+h}) - \chi]\} \\ &\quad - p_t \times \{\mathbb{E}_{t-} [B_t^{-\gamma} (F_{i,t} + (\chi - 1)B_t) - \chi]\} \\ &= (p_{t+h} - p_t) \{\mathbb{E}_{t-} [B_t^{-\gamma} (F_{i,t} + (\chi - 1)B_t) - \chi]\} \\ &= (p_{t+h} - p_t) \{\mathbb{E}_{t-} [B_t^{-\gamma} F_{i,t}] + (\chi - 1)\mathbb{E}_{t-} [B_t^{1-\gamma}] - \chi\} \\ &\quad (p_{t+h} - p_t) \{\mathbb{E}_{t-} [B_t^{-\gamma}] \mathbb{E}_{t-} [F_{i,t}] + (\chi - 1)\mathbb{E}_{t-} [B_t^{1-\gamma}] - \chi\} \end{aligned} \quad (50)$$

where the second line comes from the assumption that expectations of F and B are constant over short horizons (in this case $h = 1/12$), and the third line comes from the assumption that F and B are independent.

Plugging in Equation (50) into Equation (49), we get:

$$\log\left(\frac{P_{i,t+h}}{P_{i,t}}\right) = \log\left(\frac{D_{i,t+h}}{D_{i,t}}\right) + (p_{t+h} - p_t) \underbrace{\frac{\{\mathbb{E}_{t-} [B_t^{-\gamma}] \mathbb{E}_{t-} [F_{i,t}] + (\chi - 1)\mathbb{E}_{t-} [B_t^{1-\gamma}] - \chi\}}{\delta_i + \phi_{H,i}}}_{\equiv \tilde{\beta}_i} \quad (51)$$

The preceding equation is a regression of returns on dividend growth and changes in p_t . The regression coefficient $\tilde{\beta}_i$ is an increasing function of $\mathbb{E}_{t-} [F_{i,t}]$ since $\mathbb{E}_{t-} [B_t^{-\gamma}] > 0$, $\delta_i > 0$, and $\phi_{H,i} > 0$.

In practice, the explanatory variable in the regression is $p_{t+h}^* - p_t^*$, as opposed to $p_{t+h} - p_t$. As in the main text, I use β_i to denote the regression coefficient when using the risk-neutral probability of disaster. So how does β_i differ from $\tilde{\beta}_i$ in terms of its functional dependence on $\mathbb{E}_{t-} [F_{i,t}]$? Thus, because the mapping between p_t and p_t^* has nothing to do with $\mathbb{E}_{t-} [F_{i,t}]$, β_i will still be increasing in $\mathbb{E}_{t-} [F_{i,t}]$.

Inherently, this proof rests on the assumption that when estimate β_i , the expectations of B and F stay constant over the estimation horizon. When I implement the procedure empirically, I use trailing

one-year data to estimate the regression, so this assumption is likely to hold true.

C.8 Proof of Corollary 1

The proof is an extension of Equation (51). First, notice that with a constant B and F :

$$\log\left(\frac{P_{i,t+h}}{P_{i,t}}\right) = \log\left(\frac{D_{i,t+h}}{D_{i,t}}\right) + (p_{t+h} - p_t) \times \frac{B^{-\gamma}(F_i + (\chi - 1)B) - \chi}{\delta_i + \phi_{H,i}} \quad (52)$$

Next, recall that our estimate of p_t^* is multiplied by a scalar $\bar{\Psi}$, which represents the average Ψ_i in the economy. Recall from the main text that $\Psi_i = 2 \times (1 + \ln(F_i) + [\ln(F_i)]^2 / 2 - F_i)$. Thus, my estimate of $\hat{p}_t^* = \bar{\Psi} p_t^*$.

Under the assumption that p_t is constant within each month (τ), it follows that:

$$\begin{aligned} \widehat{p_{t+h}^*} - \widehat{p_t^*} &= \bar{\Psi}(p_{t+h}^* - p_t^*) / \tau \\ &= \bar{\Psi} B^{-\gamma} \left(\int_{t+h}^{t+h+\tau} p_s ds - \int_t^{t+\tau} p_s ds \right) / \tau \\ &= \bar{\Psi} B^{-\gamma} (p_{t+h} - p_t) \\ &\Leftrightarrow \\ p_{t+h} - p_t &= \frac{\widehat{p_{t+h}^*} - \widehat{p_t^*}}{\bar{\Psi} B^{-\gamma}} \end{aligned} \quad (53)$$

Substituting Equation (53) into (52) yields the desired result:

$$\begin{aligned} \log\left(\frac{P_{i,t+h}}{P_{i,t}}\right) &= \log\left(\frac{D_{i,t+h}}{D_{i,t}}\right) + \left(\widehat{p_{t+h}^*} - \widehat{p_t^*}\right) \times \frac{B^{-\gamma}(F_i + (\chi - 1)B) - \chi}{\theta B^{-\gamma}(\delta_i + \phi_{H,i})} \\ &= \log\left(\frac{D_{i,t+h}}{D_{i,t}}\right) + \left(\widehat{p_{t+h}^*} - \widehat{p_t^*}\right) \times \frac{B^{-\gamma}(F_i + (\chi - 1)B) - \chi}{2 \times \{1 + \ln(F_{SP}) + [\ln(F_{SP})]^2 / 2 - F_{SP}\} \times B^{-\gamma}(\delta_i + \phi_{H,i})} \end{aligned}$$

C.9 How Big Are the “Higher Order” Terms for Resilience?

In both the definition of firm-level resilience and the resilience of the consumption claim, the “higher order” terms $v^2(k_t)$ appears frequently. If we can ignore this term, then the expressions for the resilience of the consumption claim is given by $H_{ct} = k_t - \gamma(1 - \gamma) \frac{\sigma_c^2}{2}$. For the case of the consumption claim, this means that the price-dividend ratio will be exactly linear in H_{ct} (as opposed to k_t , see expression (24)). In his Online Appendix, [Gabaix \(2012\)](#) shows $v(k_t)$ is proportional to the volatility of p_t , which he calibrates to $\sigma_p = 0.4\%$. In this paper, the only difference between my definition of consumption resilience in (22) from his same expression is a constant; thus, the analysis of the magnitude $v(k_t)$ is the same, since this determines the volatility of resilience. [Gabaix \(2012\)](#) finds that, for example, the $v(k_t)$ terms contribute less than 1% to quantities such as H_{ct} and H_{it} , and for the same reason, the approximations used in the main text are a very accurate version of the fully solved model presented in this appendix.