The Pricing of Disaster Risk*

Emil Siriwardane†

NYU Stern School of Business

First Draft: Sept. 2011
Current Draft: May 9, 2013

Abstract

I analyze a model economy with rare disasters that yields a theoretically-grounded measure of firm disaster risk that can be extracted from option prices. Specifically, I develop a simple option strategy that reflects only the risk inherent in the disaster states of the economy. My disaster risk measure (DR) proxies for the ex-ante disaster risk of a firm's stock, thereby circumventing the difficult task of using historical equity returns to estimate disaster exposure. In equilibrium, firms with high disaster risk naturally require higher returns on average. Indeed a zero-cost equity portfolio that is exposed to high disaster risk stocks earns excess annualized returns of 12.13%, even after controlling for standard Fama-French, momentum, liquidity, and volatility risk-factors. Moreover, the model suggests how to use my DR measure to infer the risk-neutral probability of a consumption disaster. Cross-sectionally, assets with higher DR demonstrate higher price sensitivity to changes in the probability of a consumption disaster. These results are also fully consistent with the model since these assets carry the most disaster risk.

*I am especially thankful to Robert Engle and Xavier Gabaix for their advice on this research. In addition, I am grateful to Viral Acharya, Yakov Amihud, Jennifer Carpenter, Itamar Drechsler, Eric Ghysels, Paul Glasserman, Nikunj Kapadia, Bryan Kelly, Stijn Van Nieuwerburgh, Alexi Savov, and Josie Smith for many valuable discussions and comments. I also thank seminar participants at NYU Stern, the MFM Spring 2013 Meetings, and the Office of Financial Research. The author wishes to acknowledge grant support from the MFM and the Center for Global Economy and Business at NYU Stern.

†Affiliation: NYU Stern School of Business. Address: 44 West 4th St., Floor 9-Room 197K. New York, NY 10012. E-mail:esiriwar@stern.nyu.edu.

1 Introduction

There has been a surge in disaster models and tail risk measures in the wake of the Great Recession of 2008. From an asset pricing perspective, a natural implication of disaster models is stocks that are exposed to disaster risks should earn higher average returns. Empirically, the difficulty in producing a viable measure of a stock's disaster risk is that disasters are by definition rarely realized. Thus, measuring how a firm performs when a latent disaster process strikes remains a challenge from an econometric perspective. Nonetheless, pinning down the compensation for bearing disaster risk necessitates overcoming this measurement issue, which is exactly what I set out to do in this paper.

I start by proposing a simple measure of a stock's disaster risk that is theoretically-grounded in a model economy that experiences consumption disasters with a time-varying probability. My approach overcomes the difficulty of using historical time series (or consumption) data by recognizing that if performance during a potential disaster is indeed priced into the equity return of a firm, it will must also be priced into options on that same firm. This insight allows me to use *options* to back out the disaster risk of a large number of firms, which I then use to form a zero-cost *equity* portfolio that is exposed to firms with high disaster risk. Options in this setting are thus a means to an end for selecting stocks whose option prices imply high covariance with the latent disaster variable that affects the prevailing stochastic discount factor in the economy. I call the equity portfolio that is formed by sorting on my disaster risk (DR) measure H_{DR}, where the H means "high disaster risk".

A natural starting point for measuring a stock's disaster risk is simply by looking at the price of out of the money (OTM) put options on the firm; however, there are subtle problems with using simple OTM put options to measure a firm's disaster risk. Empirically, deep OTM put options are thinly traded, suffer from liquidity biases, and in some cases may even suffer price distortions from implicit government guarantees.[1] Still, if one restricts attention to put options that are not extremely deep OTM, the value of the put option will be comprised of risk that occurs in "normal times" and also disaster risk. I solve this problem by showing that if in normal times the distribution of returns is roughly symmetric (over the relevant horizon) then a symmetrical OTM call option will capture the same normal-time value that its put option counterpart contains, yet will not contain any value derived from disaster risk.[2] Therefore, one can subtract out a symmetrical call option from a put option to isolate a DR measure for the stock. I use an economic model to relate this DR measure to expectations about how the *stock* will covary with the stochastic discount factor. To this end, the economic disaster model allows me to relate the DR measure backed out from options on a firm to the portion of that firm's equity premium that comes from bearing disaster risk. The model naturally predicts firms with a higher DR measure should have a higher equity premium, which is the basis for forming the H_{DR} equity portfolio in the first place.

Empirically, I find that the H_{DR} portfolio earns excess annualized returns of 12.13% after controlling for standard risk-factors like the Fama-French factors, momentum, Pastor-Stambaugh liquidity, and aggregate volatility. In constructing a DR measure for a given firm, I am careful to control for changes in volatility since I am interested in focusing on disaster risk and not volatility risk. The H_{DR} portfolio has a very low loading on innovations to aggregate volatility, which confirms that the large excess returns I observe are not being driven by exposure to volatility.

[1] In addition, recent research by Kelly, Lustig, and Van Nieuwerburgh (2012) suggests that OTM put option prices may be distorted by implicit government guarantees.

[2] In unreported theoretical results, I relax the assumption of asymmetry in the model by allowing "idiosyncratic" disasters to occur in normal times. As long as the probability and severity of idiosyncratic disasters is not too heterogenous across firms, then my DR measure will still serve as a valid proxy for covariance with the latent disaster process.

Since my DR measure and its relationship to equity premiums derives (in part) from a disaster model, I conduct some additional empirical tests of the model in order to validate the disaster content of my DR measure. The model suggests a theoretical way to back out the risk-neutral probability of a disaster using linear combinations of my DR measure. Moreover, since the hidden disaster process affects all firms, the probability of a disaster can be inferred from *any* firm-level measure of DR; hence, I back out the risk-neutral probability of disaster based on individual firm DR measures. Using principal component analysis, I find a large principal component common to the panel of disaster probabilities extracted from firm DR measures, with the first principal component capturing nearly 66% of all variation. The factor extracted from firm implied disaster probabilities is able to forecast macroeconomic time-series such as changes in unemployment and growth in industrial production, which confirms that the probability of disaster extracted from the panel of individual firms is working at a macroeconomic level.

The disaster model then suggests a precise contemporaneous relationship between realized returns on the market and changes in the risk-neutral probability of a disaster. Intuitively, when the probability of a disaster increases, the market should fall in price. The model's calibration of market price sensitivity to change in disaster probabilities indeed matches the empirical estimates quite well, thereby providing additional evidence for the disaster model and the content of my DR measure.

Additionally, I explore a cross-sectional test of the model/DR-measure. The disaster model says that firms with higher DR measures should experience larger price drops when the probability of disaster increases. Using the probability of disaster extracted from firm implied disaster probabilities, I run a simple regression of realized *firm* level returns on contemporaneous changes in the probability of disaster. I empirically confirm that firms with a high DR measure indeed have a larger negative contemporaneous relationship with changes in disaster probability. Taken together, the three empirical tests of the disaster model serve as validation of my DR measure and its use in forming the HML_{DR} portfolio. Moreover, I view these tests as providing strong theoretical support for why H_{DR} captures the premium for bearing aggregate disaster risk.

The remainder of the paper is follows. In Section 2, I briefly discuss other work that has been conducted on disaster risk measurement and its effect on asset prices. Section 3 quickly introduces my DR measure, and then uses this measure to form equity portfolios. The latter part of this section goes through the standard empirical asset pricing tests of these portfolios. Next, in Section 4 I provide the economic underpinnings to my disaster measure through a model economy that experiences rare but severe consumption disasters. It is here that I will spend considerable time developing the asset pricing implications of equities and options for a given firm in a disaster environment. Section 5 gives a brief overview of the data used in this study, and Section 6 explores additional empirical implications of my disaster measure and the model economy. Finally, Section 7 concludes with a discussion of my measure in relation to other measures of downside risk and some areas of future research.

2 Related Literature

2.1 Tail Risk and Asset Prices

The relationship between tail risks and asset prices was first explored by Rietz (1988) as a solution to the equity premium puzzle. Rietz extends the Mehra-Prescott (1985) specification of the aggregate consumption growth process to include a rare crash state. By doing so, he is able to simultaneously generate a high enough equity risk premium and a low enough risk-free rate, while still using reasonable assumptions on risk aversion in the economy. It is important to recognize that the key mechanism that was borne out of this literature was that a *consumption* disaster can generate meaningful effects on equity premiums by affecting the stochastic discount factor with which assets are valued in equilibrium. The long standing criticism of this paper focused on its estimates of the probability and magnitude of a rare disaster (tail event), which were deemed by most in the field to be unreasonable given historical U.S. consumption growth.

Barro (2006) revives this formulation by extending the Rietz (1988) model and calibrating the magnitude and probability of economic disasters to match historical international data. Barro's (2006) calibration then delivers reasonable estimates for the equity premium, among other asset puzzles he considers in the paper. Gabaix (2012) extends the Barro-Rietz framework to accommodate time-varying intensity of disasters and provided closed-form solutions to a number of asset pricing puzzles. Wachter (2012) offers a variant of the Gabaix (2012) model by considering the effects of a time-varying probability of a consumption disaster (albeit a constant severity of disaster) in an economy with recursive preferences instead of power utility. The use of recursive preferences enables Wachter (2012) to reconcile some of the implications from the calibration portion of Gabaix (2012), namely the behavior of the risk-free rate. I will discuss the frameworks of these two models in a more detailed fashion in Section 4 when I lay out my modeling approach. In all of these models, assets that payoff during times of high tail risk command a lower equity premium over a risk-free asset.

A parallel literature evolved within the framework of the long run risks model of Bansal and Yaron (2004). In their model, the prevalent measure of risk in the economy is covariance with long-run consumption growth. Equity premiums are shown to have two components of risk: consumption growth itself and consumption volatility. Kelly (2011) extends this model by subjecting log-consumption growth to 2 shocks, a normal and a heavy tailed shock. The result is that equity premiums have three sources of risk, with tail risk additionally forecasting returns in equilibrium. The tail risk process that affects both individual firms and aggregate consumption growth also implied an estimation technique for a tail risk factor, which will be further discussed below.

2.2 Estimating Tail Risk

As previously discussed, Barro (2006) estimates the parameters of his tail risk process using aggregate consumption growth data on an international level. Kelly (2011) employs a novel estimation technique motivated by extreme value theory. In particular, he shows that because individual firm tail risk is closely related to aggregate tail risk then in a large enough cross section a non-trivial number of firms will experience tail events. Furthermore, he argues that the tail of any given asset follows a power law. This motivates the use of Hill's (1975) power law estimator to a cross section of stocks over time, thereby providing a time-varying measure of tail risk. A very appealing feature of this estimation technique is the ability to use a large panel of equities to back out an aggregate tail risk factor. One large difference between the approach I take in this paper and Kelly (2011) is that his approach will inherently be backward looking since it relies on historical return data. In fact, any econometric measure that relies on historical equity data will suffer from the criticism that historical equity returns may not contain enough disaster information given the scarcity of disasters in the time-series.

Bollerslev and Todorov (2011) use short-dated OTM options on the S&P 500 to estimate a risk-neutral tail measure. They take advantage of the notion that short-maturity OTM options have little value unless a rare event occurs before expiration. They construct a nonparametric estimator of the right-tail from call options and of the left-tail from put options.[3] Du and Kapadia (2012) also develop an option-based model free measure of tail risk. This approach capitalizes on the idea that the difference between quadratic variation and integrated variance should isolate the risk-neutral jump intensity in a general class of jump-diffusion models. In fact, Du and Kapadia's (2012) measure highly resembles the disaster risk measure developed in this paper; however, my approach not only enables me to comment on why my measure captures disaster risk economically, but also why disaster risk is related to equity risk premiums. Backus, Chernov, and Martin (2011) use a version of the Barro-Rietz hypothesis that models consumption growth as having an independent normal component and an independent Poisson mixture of normals to capture jumps. They then specify preferences in their economy to derive the well known link between the physical distribution, the pricing kernel, and the risk-neutral distribution.[4] From options on the S&P 500, they estimated the parameters of the risk-neutral distribution, then used a power utility function to back out the parameters of the physical distribution of equity returns.[5] One significant difference in my approach is that I allow for a time-varying probability of disaster, whereas Broadie, Chernov, and Johannes use a constant probability.

Perhaps the closest related methodology to mine is Farhi et al. (2013), who use what they call risk-reversals (I call this disaster risk) from currency options to isolate relative disaster risk between countries. Their main object of study is the excess return of currency carry trades, whereas I am interested in how rare disasters affect stocks. I thus diverge from their approach by developing disaster risk as a stock-specific characteristic and explore the

[3] In fact, Bollerslev and Todorov (2011) define their "fear index" as the difference between the portion of the variance risk premium coming from negative jumps and positive jumps. Practically speaking, this will resemble a my DR measure, though the authors do not make this point explicitly.

[4] Note in their paper these are the relevant distributions for equity returns. There is an additional assumption made that models equity as a claim to a portion of aggregate consumption, which completes the connection between the equity return process and the consumption growth process

[5] Estimation of the risk-neutral distribution parameters follows an earlier paper by Broadie, Chernov, and Johannes (2007)

relationship between disaster risk and stock returns.

Many of these measures are designed to be "model-free", and are derived using basic assumptions on the return generating process. These measures largely focus on the econometric issues with estimating jump risks inherent in a stock price. This paper can be viewed as a simpler counterpart to the sophisticated econometric techniques presented on the literature since the estimation of tail risk is motivated from an economic model instead of relying on extreme value theory or assumptions on the underlying return process. The advantage of the approach taken in the literature in estimating tail risk is that by writing down a model of the risk-neutral return process that includes jumps, the resulting estimation technique is then internally consistent and can be easily interpreted as estimating a jump measure. My approach is instead is internally consistent with the economic model put forth in this paper. Furthermore, I am able to assess the economic significance of my disaster risk measure and link it to asset prices through an economic model.

3 A Disaster Risk Equity Portfolio

The bird's eye view of disaster risk is simple: if investors care about how a stock will perform in a disaster, then stocks that are expected to perform poorly in a disaster will require a premium in equilibrium. A straightforward implication of this idea is that for a viable measure of a firm's disaster risk, stocks with high disaster risk will, on average, earn higher returns than stocks with low disaster risk. The first challenge in quantifying the price of disaster risk in stock returns (if any) is then producing a measure of disaster risk. To this end, I will use *options* to tell me which firms have high disaster risk. Historical time series of stock returns is a difficult source of quantifying disaster risk since disasters are rare and hard to identify. On the other hand, options are forward looking and naturally will have expectations about a firm's performance built into their prices. Intuitively, the same expectations about a firm's disaster performance that determine option prices will also play a role in determining the premium required for holding the equity of the firm. For now, I will take as given that my disaster risk (DR) measure accurately reflects a stock's expected performance in a disaster and proceed with standard empirical tests for whether disaster risk is priced. Then in section 4, I will use a model economy that experience consumption disasters to show that my (DR) measure is a theoretically-grounded measure of disaster risk for a firm's equity.

3.1 Forming Equity Portfolios Based on a Disaster-Risk Measure

First, define $Put_t^i(K;\tau)$ as the value at time t of a put option on firm i at strike K with maturity τ. Embedded in this definition are the standard underlying parameters such as spot price of the firm's stock, the prevailing risk-free interest rate, and the volatility of the firm over the lifetime of the option. I suppress the functional dependence on standard option parameters for now since these are not important in developing the intuition of my measure. The call pricing function is defined analogously as $Call_t^i(K;\tau)$. For a given firm i, I construct a measure of disaster risk on date t using options on firm i as follows:

$$
\begin{aligned}
DR_t^i &\equiv \frac{Put_t^i(m \cdot S_t^i;\tau) - m \times Call_t^i(m^{-1} \cdot S_t^i;\tau)}{S_t^i} \\
m &= K/S_t^i
\end{aligned} \tag{1}
$$

S_t^i is the stock price for firm i on date t and m is the moneyness of the put option. m will always be less than 1, which implies that both the put and the call option will be out of the money (OTM). One advantage of my DR measure is its simplicity: it amounts to the price of an OTM put option minus a symmetrical OTM call option, appropriately normalized by the firm's current stock price. Normalizing by the stock price S_t^i allows me to compare my DR_t^i measure across firms. Formally, the normalization means the options are being evaluated in terms of returns instead of price levels. As mentioned, I will spend considerable time motivating my DR_t^i measure in section 4, but it is worth providing some basic intuition at this juncture. Part of a firm's put option value comes from what occurs in "normal" times and what occurs in a disaster. The call option also has value from what occurs in normal times; however, the call option has no value from what happens in a disaster since prices will fall. If in normal times the returns of a firm are (roughly) symmetric then the normal-time value of the put will equal the normal-time value of the symmetric call. Hence, subtracting the symmetrical call cancels out the normal-time value of the put option, leaving only the portion of the put value that is due to disaster risk. Firms with high DR_t^i are thus expected to perform worse in a disaster, and if this risk is priced, then these firms will require higher equity returns in equilibrium. Indeed, it may be the case that this measure of disaster risk picks up only "idiosyncratic" disasters. If this were the case though, then standard portfolio theory would say that stocks with high "idiosyncratic" disaster risk would not earn a premium on average. As we will see shortly, this is not the case.

For the remainder of the paper, I set $\tau = 1$ month. The big issue in constructing DR_t^i is then choosing the appropriate moneyness for the put option. In order to control for the volatility of firm i, it is natural to let m be a function of firm specific volatility. An easy function that maps firm implied volatility to moneyness is the Black-Scholes (BS) delta. Hence, I use put and call options such that the BS delta is 25 (in absolute value). In section 5, I provide the full details and motivations for constructing DR_t^i. The important implication of setting m based on firm implied volatility is that comparing DR measures across firms will not necessarily be an exercise in comparing implied volatilities.[6]

The next step in honing down whether disaster risk is priced is forming an equity portfolio based on my DR_t^i measure. To do this, I construct a daily DR_t^i measure for all firms in the S&P 500 going back to 1996. To avoid any lookahead bias, I determine the constituents of the S&P 500 at the end of each month with data downloaded from Compustat. Then, I form an *equity* portfolio based on the DR_t^i measure for all firms over the previous month. That is, at the end of each month I sort the stocks in the S&P 500 into quintiles based on their median DR_t^i measure over the previous month. Since my DR measure comes from the value of put and call options, this is the sense in which firm-level options are a means to an end for forming an equity portfolio.[7] Henceforth, I will refer to the equity portfolios formed based on this procedure as the disaster risk portfolios.

3.2 The Price of Disaster Risk

Table 1 contains basic summary statistics for the excess returns of the 5 disaster risk portfolios. A striking feature of these portfolios is the pattern of average returns when moving from low disaster risk to high disaster risk. Intuitively, average returns and volatility are roughly increasing in disaster risk; however, the Sharpe ratios of the disaster risk portfolios stays relatively flat. As one might expect, the minimum monthly return on the higher disaster

[6] The reader may notice that DR_t^i is closely related to skewness. However, by setting m to be a function of implied volatility, my DR_t^i measure will be distinct from skewness precisely because cross-sectional heterogeneity in DR_t^i is not necessarily due to heterogeneity in implied volatilities. A typical skewness measure will not share this feature, but the two measures will still be closely related. Still, I show theoretically why sorting stocks on my DR measure generates a premium and thus view skewness as reflecting aspects of my DR measure, not vice versa.

[7] The weights in this portfolio are a function of DR_{it}. In the high quintile portfolios I weight according to disaster risk and in the lower decile portfolios I weight according to the inverse of disaster risk. Full details are in the Online Appendix.

risk portfolios is quite large around -25% (these all occurred in November 2008). The higher disaster risk portfolios reward investors with extremely high returns in some months. In the high disaster risk portfolio, the maximum return is as high as 44.83% in a given month.

Table 1: Summary Statistics for Equity Portfolios Formed on Disaster Risk

Disaster Quintile	**Average Return (%)**	**Volatility (%)**	**Sharpe Ratio**	**Min (%)**	**Max (%)**
Low Disaster Risk	6.32	12.93	0.49	-15.16	11.86
2	8.02	15.20	0.53	-18.36	13.81
3	6.54	17.56	0.37	-20.03	18.49
4	10.24	21.12	0.49	-24.69	22.72
High Disaster Risk	16.43	33.95	0.48	-25.07	44.83

Notes: The table reports summary statistics of the excess returns of disaster risk quintile portfolios formed based upon my DR measure. At the end of each month m, I stocks in the S&P 500 into terciles based on their median value of DR_i over the month m. I then hold the equity portfolio from the end of month m to the end of month $m+1$, then rebalance. The first three columns present annualized values. The minimum and maximum values are for *monthly* observations. The sample period is from January 1996 to October 2013.

To get a better feel for the time-series of the disaster portfolios, I plot the low DR portfolio and the high DR portfolio in Figure 1. The extreme returns of the high disaster risk portfolio are quite striking, so much so that I investigated some of the more extreme months by hand to make sure there was not something wrong with the data. The extreme returns of the high DR portfolio are in fact accurate. For example, the large return in April of 2009 came from purchasing stocks such as Ford, Principle Financial Group, American Express, and Wynn Resorts. All of these stocks experienced nearly 100% returns in this month, but also ranked among the stocks with the highest DR in March 2009.[8] Indeed there are months where the high DR portfolio delivers enormous returns, but there are also quite a few months where this portfolio loses more than 20%. These months coincide with some well known "crisis" periods

[8]A complete breakdown of the stocks in each DR portfolio over time is available upon email request.

such as the LTCM fallout and the September 11, 2001 terrorist attacks on the United States, which I took as good news since this portfolio should be exposed to these types of events. To better access the nature of the risks in these disaster portfolios, I turn to time-series regressions of excess returns on common risk-factors.

Figure 1: Equity Portfolios Formed Based on Disaster Risk

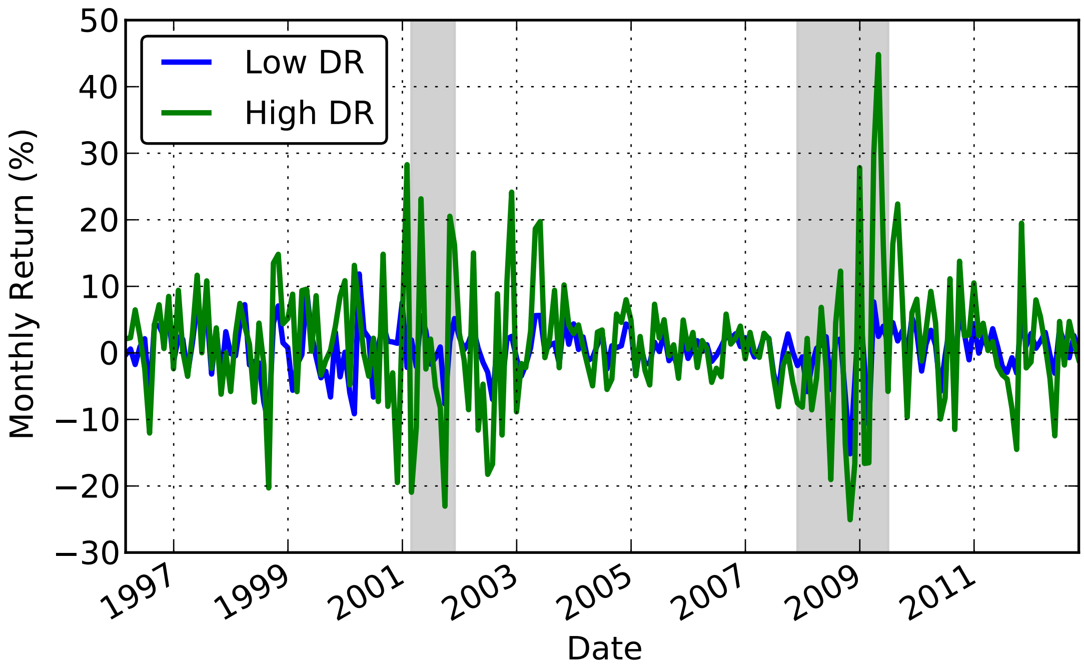

Notes: This figure plots monthly excess returns of three different equity portfolios formed based on disaster risk. The blue line represents low disaster risk, and the green line is the high DR portfolio. Shaded regions indicated NBER recession dates.

Table 2 contains time-series regressions of the disaster portfolio excess returns on the Fama-French (1993) factors, momentum, the Pastor-Stambaugh (2003) traded liquidity factor, and the return on the VIX.[9] The pricing errors (α's) and the loadings on the risk-factors

[9]To calculate the return on the VIX I simply treated the VIX as if it were a stock. Using changes in the VIX does not change the results in any way, as the two series are highly correlated. In practice, the VIX is

Table 2: Risk-Adjusted Returns for Equity Portfolios Formed on Disaster Risk

Disaster Quintile	**Annualized α**	$\boldsymbol{\beta_{mkt}}$	$\boldsymbol{\beta_{smb}}$	$\boldsymbol{\beta_{hml}}$	$\boldsymbol{\beta_{mom}}$	$\boldsymbol{\beta_{liq}}$	$\boldsymbol{\beta_{vol}}$
Low Disaster Risk	1.97	0.47	0.10	0.28	-0.10	0.07	-0.06
	(1.11)	(7.10)	(1.78)	(2.96)	(-1.07)	(1.82)	(-2.04)
2	2.73	0.64	0.14	0.31	-0.12	0.07	-0.06
	(1.79)	(11.30)	(2.09)	(3.51)	(-1.61)	(1.94)	(-2.23)
3	0.58	0.79	0.12	0.31	-0.17	0.07	-0.05
	(0.34)	(15.08)	(2.30)	(3.44)	(-2.22)	(2.40)	(-2.00)
4	3.38	0.98	0.25	0.37	-0.21	0.06	-0.05
	(1.59)	(20.60)	(4.77)	(4.61)	(-4.49)	(1.38)	(-1.59)
High Disaster Risk	12.13	1.46	0.22	0.13	-0.54	-0.20	-0.07
	(2.51)	(13.20)	(1.51)	(1.33)	(-5.03)	(-2.30)	(-1.50)
HML_{DR}	10.17	0.99	0.12	-0.14	-0.44	-0.27	-0.02
	(1.76)	(6.25)	(0.71)	(-0.79)	(-2.28)	(-2.44)	(-0.36)

Notes: The table reports estimated coefficients from the time-series regression $R_i - R_t^f = \alpha_i + \beta_{i,m}(R_{mkt,t} - R_t^f) + \beta_{i,s}SMB_t + \beta_{i,h}HML_t + \beta_{i,mom}MOM + \beta_{i,liq}PS_{liq} + \beta_{i,vol}RVIX + \varepsilon_{i,t}$. The sample period is from January 1996 to October 2013. All t-statistics are calculated using Newey-West (1987) HAC standard errors are listed below the point estimates in parenthesis. The 3 Fama-French were taken from Ken French's website, as well as the momentum factor. The factor PS_{liq} corresponds to the Pastor-Stambaugh (2003) traded liquidity factor, and was taken directly from Robert Stambaugh's website. Finally, $RVIX$ is simply the monthly excess "log-return" on the VIX closing prices.

were calculated using GMM on the entire system of portfolios. The covariance matrix for the estimates accounts for contemporaneous correlation in the pricing errors across portfolios, as well as heteroskedasticity and autocorrelation within each portfolio. The first thing to notice is the α of the high disaster risk portfolio is quite large at 12.13%. Indeed, the market beta (β_m) is increases when moving from low to high disaster risk, but not enough to account for the excess returns the pattern of excess returns observed in the disaster portfolios. The high

itself not a traded volatility factor. Ang et. al (2006) also construct a factor mimicking portfolio $FVIX$ for the VIX. They argue that at a monthly frequency using a factor mimicking portfolio may be more important since the conditional mean of the VIX is non-negligible. The authors look at first differences in the VIX rather than treating the VIX as a stock. To start I assume that the return on the VIX will serve as a useful enough proxy to filter out volatility exposures, but will conduct robustness checks to make sure.

disaster risk portfolios have especially high market betas, which may be alarming at first when interpreting the α of this portfolio as disaster compensation. Still, I show theoretically in the Online Appendix that a high disaster risk portfolio will likely have a high market beta, even in an economy where exposure to the market is not even priced. The reason is that both the market and the high disaster risk portfolios will have high exposures to a latent disaster variable, and thus the high disaster risk portfolio will have high covariance with the market.[10]

Interestingly, all of the disaster portfolios except the high disaster risk quintile have a positive and significant loading on HML. In fact, the load on HML drops substantially in the high disaster risk portfolio, which is a bit surprising. The value premium has been linked to distress costs so one would expect the high disaster portfolio to have the highest loading on HML. Turning to the SMB factor, the loadings on the SMB factor increase when moving from low disaster risk to high disaster risk, but is not statistically significant at a 5% confidence level for the highest quintile. These results suggest that small stocks may more exposed to disaster risk. In light of the recent crisis in which small businesses have been hit hard by the credit crunch, this finding is quite intuitive. The combined loadings on the SMB and HML factors suggest there might be some hope in using disaster risk to explain the cross-section of the Fama-French portfolios that are sorted across size and book to market ratios. I will explore a potential link between disaster risk and the value and size premiums further in the Online Appendix.

Importantly, letting the moneyness in my option-based DR measure vary with volatility was successful in ensuring the equity disaster portfolios are immune to aggregate volatility.

[10] Also in the Online Appendix, I do a sanity check to make sure that value-weighting these portfolios doesn't earn alpha. This is because a value-weighted sum is essentially the market itself. Indeed, this is the case.

To see why, notice that the loadings β_{vol} on returns (changes) to the VIX are close to zero for all disaster portfolios. Moreover, if volatility were driving these results we would expect to see high disaster risk portfolios have larger loadings on volatility innovations, but clearly the loadings across all portfolios are basically the same. My proxy for innovations to aggregate volatility in this case was the monthly growth in the VIX. Since this may be a poor proxy, in the next subsection I ensure that the large excess returns generated by the disaster portfolios are robust to other measures of aggregate volatility exposure.

In addition, there is a monotonic pattern in the negative exposure of the disaster portfolios to a momentum strategy, with the momentum loading of the higher DR portfolio also having statistical significance. The last risk-factor I included in the time-series regression was the traded liquidity factor of Pastor and Stambaugh (2003). All but the high DR portfolio have a positive and significant loading on the liquidity factor. However, the high disaster risk portfolio have a slightly negative loading on innovations to aggregate liquidity. This finding is not surprising though, as high disaster risk tends to associate with low levels of liquidity.

Additionally, I conduct a joint test that all of the α's of the disaster portfolios are zero using standard GMM asymptotic distribution theory. The value of the sample statistic is 18.2 and is $\chi(5)$ distributed, indicating a strong rejection of the null that the disaster portfolio α's are jointly zero. The economic significance of 12.13% in risk-adjusted returns for the high DR portfolio is powerful. In a sense, I chose a subset of stocks that should be hardest to "earn alpha" given these risk-factors (i.e. stocks in the S&P 500). The fact that I am able to do so with such a large amount is, in my opinion, confirmation that disaster risk matters. These results together suggest that these disaster portfolios are exposed to additional risks not captured by the F-F model, momentum, aggregate liquidity, and volatility risk-factors. Since the portfolios are constructed based on my DR measure, I take this as strong evidence of a premium for bearing disaster risk. Finally, because the high disaster risk portfolio is zero-cost, I quantify the compensation for bearing disaster risk as 12.13% in annualized returns.

3.3 Robustness to Aggregate Volatility Exposure

The control for exposure to aggregate volatility in Section 3.2 was essentially using monthly changes in the VIX as an aggregate volatility factor. However, as Ang et. al (2006) point out, using the monthly change in VIX is a poor approximation for innovations in aggregate volatility. In order to ensure that the excess returns earned by the disaster portfolios is not in fact compensation for aggregate volatility risk, I consider an alternative test in this section. The strategy I employ mirrors Ang et. al. (2006), which I now outline in detail:

1. Index the daily return observations within month t for firm i as $d = 1, ..., t_i$. At the end of each month t, run the following regression for each firm i:

$$r_{id} = a + \beta^i_{VIX,t} \times \Delta VIX_d + \varepsilon_{i,d}, \qquad d = 1, ..., t_i$$

   In other words, within each month regress the *daily* returns for each firm on the *daily* changes in the VIX. I will call $\beta^i_{VIX,t}$ the "VIX-beta" for firm i in month t, since it is a measure of how much a stock's return covaries with innovations to aggregate volatility. Daily changes in the VIX are a good proxy for innovations to aggregate volatility since the VIX is highly autocorrelated at daily frequencies. This approach is therefore meant to overcome the potential issue with using monthly changes in the VIX as a proxy for innovations to aggregate volatility.

2. Run a cross-sectional regression of disaster risk on the VIX-beta for each firm:

$$DR_{i,t} = c_t + \gamma_t \beta^i_{VIX,t} + \xi_{i,t}$$

3. Form monthly equity portfolios based on the residuals $\xi_{i,t}$. Firms with the highest $\xi_{i,t}$ go into the high disaster risk bucket, an so forth. If the equity premium from bearing disaster risk is in fact being driven by exposure to aggregate volatility, then portfolios sorted on $\xi_{i,t}$ should earn substantially less excess returns than sorts based on $DR_{i,t}$ alone.

Table 3: PRICE OF DISASTER RISK, CONTROLLING FOR AGGREGATE VOLATILITY

Quintile	(a) % α sorted on DR_{it}	(b) % α sorted on ξ_{it}
Low	1.66	3.73
	(0.93)	*(1.63)*
2	2.40	1.05
	(1.55)	*(0.68)*
3	0.28	1.07
	(0.16)	*(0.71)*
4	3.10	1.37
	(1.43)	*(0.79)*
High	11.73	12.46
	(2.40)	*(2.86)*
HML_{DR}	10.06	8.73
	(1.76)	*(2.16)*
χ^2	18.2	8.93
	$p = 0.003$	p=0.11

Notes: The table reports excess annualized portfolio α after controlling for standard Fama-French factors, momentum, and Pastor-Stambaugh liquidity. In column (a) equity portfolios are formed by sorting on DR_{it}. In column (b), equity portfolios are formed by sorting on ξ_{it}, where ξ_{it} are the residuals from monthly regressions of DR_{it} on $\beta^i_{VIX,t}$. T-statistics are listed in paretheses, and are calculated using GMM with HAC standard errors. The last row of the table is the test that all α are jointly zero. Data spans January 1996 to January 2012.

Table 3 presents the results from the original sorts based on DR_{it} alone and also after controlling for exposure to aggregate volatility. We can see from the results that controlling for volatility does indeed reduce the risk-adjusted returns earned by the disaster risk. This is most notably evident by the fact that the p-value for the test that all the α's are zero is now just above a 10% confidence threshold, whereas before it was a strong rejection. Still, the high portfolio earns 11.73% in annualized alpha before controlling for volatility, but the excess return actually increases slightly to 12.46% after controlling for volatility. Moreover, the t-statistic on the high disaster risk portfolio remains quite high at 2.86. Since the p-value for sorts based on ξ_{it} is still close to the 10% line and the excess alpha of the high disaster risk portfolio remains largely unchanged, I think it is safe to conclude that volatility can not explain the disaster risk premium.

4 A Model Economy with Rare Consumption Disasters

The model used in this paper follows Gabaix (2012), which I use to motivate my DR measure and explore additional asset pricing implications. As such, this paper spends little time developing the results in his model and instead will present only the model's most necessary components. This particular disaster model also takes advantage of linearity-generating (LG) process developed in Gabaix (2009). The reader should refer to these earlier works for a more thorough understanding of the macroeconomic dynamics in this model, though there is a short overview of LG-processes in Appendix C to guide intuition. To start, I present a tractable model economy with power utility to understand why a DR_{it} should capture properties of a disaster that are not present in simple OTM put options. I will initially skip some of the more technical details regarding equilibrium prices and realized returns in order to focus on building the intuition of my DR measure. Once I have established the intuition behind my measure, I will then pay the cost of adding Epstein-Zin preferences and provide more of the technical background for fully solving this economy. The Epstein-Zin extension is crucial for honing down the empirical predictions of the model.

4.1 Macroeconomic Environment

There is a representative agent with utility given by:

$$U = \mathbb{E}_0 \left[\sum_{t=0}^{\infty} e^{-\rho t} \frac{C_t^{1-\gamma}}{1-\gamma} \right]$$

where γ is the coefficient of relative risk aversion and $\rho > 0$ is the rate of time preferences. At each period t, the representative agent receives the consumption endowment C_t and at each period $t+1$ a disaster may happen with probability p_t. The endowment grows according to the following:

$$\frac{C_{t+1}}{C_t} = e^{g_C} \times \begin{cases} 1 & \text{if there is no disaster at } t+1 \\ B_{t+1} & \text{if there is a disaster at } t+1 \end{cases}$$

where g_C is the normal-times growth rate and $B_{t+1} > 0$ is a stochastic disaster shock that affects consumption growth in disasters. For example, if $B_{t+1} = 0.8$, consumption falls by 20% compared to normal-time growth. Indeed, this model of consumption growth has been stripped down to focus on disaster risk, since the only stochastic portion to growth enters through disasters. It obviously a more realistic description of the world to include i.i.d shocks to log-consumption growth in normal times.[11] In this case, risk-premia in the model are augmented with the familiar term involving the volatility of these i.i.d shocks and the coefficient of risk-aversion. The subsequent analysis, however, does not change and thus I exclude extra factors for the sake of parsimony.

It is straightforward to work out that the pricing kernel, or the marginal utility of consumption, evolves according to:

$$\frac{M_{t+1}}{M_t} = e^{-\delta} \times \begin{cases} 1 & \text{if there is no disaster at } t+1 \\ B_{t+1}^{-\gamma} & \text{if there is a disaster at } t+1 \end{cases}$$

where $\delta = \rho + \gamma g_C$.

4.2 Setup for Stocks

A typical stock or portfolio of stocks i is a claim on a stream of dividends $(D_{it})_{t \geq 0}$, where the dividend grows according to:

$$\frac{D_{i,t+1}}{D_{it}} = e^{g_{iD}}(1 + \varepsilon_{i,t+1}^{D}) \times \begin{cases} 1 & \text{if there is no disaster at } t+1 \\ F_{i,t+1} & \text{if there is a disaster at } t+1 \end{cases}$$

where $\varepsilon_{i,t+1}^{D} > -1$ is a mean-zero shock that is independent of the disaster event and matters only for the calibration of the dividend volatility. Notice this modeling choice accommodates a wide-array of dividend processes, which is the motivation for using it in the first place. In fact, I will take advantage of this flexibility when specifying the noise process for future

[11] In fact, Barro (2006) describes such a model. The normal-time risk premia ends up being, according to his calibration, much smaller than the disaster risk premia. This reflects the famous Mehra-Prescott (1985) equity premium puzzle.

returns, with the important point being that the modeler is (relatively) free to do so in this framework. We can think of $F_{i,t+1}$ as the recovery rate of the dividend in a disaster.

A simple way to summarize how a stock performs in a disaster is through resiliency, which is defined as $H_{it} = p_t \, \mathbb{E}_t^D[B_{t+1}^{-\gamma} F_{i,t+1} - 1]$. It is straightforward to see that assets with high resilience perform better in crises; thus, assets with higher resilience should command higher prices. Additionally, an asset with a high resilience means its performance during a disaster positively covaries with marginal utility during a disaster. It turns out it is easier to model the stochastic nature of resilience rather than p_t, B_{t+1}, and $F_{i,t+1}$ individually since this allows the use of the linearity-generating process machinery. In Appendix C, I provide a short introduction to the mechanics of LG processes as they pertain to this specific model. Formally, this means resilience evolves as follows:

$$\begin{aligned} H_{it} &= H_{i*} + \hat{H}_{it} \\ \hat{H}_{i,t+1} &= \frac{1 + H_{i*}}{1 + H_{it}} e^{-\phi_H} \hat{H}_{it} + \varepsilon_{i,t+1}^H \\ &\approx e^{-\phi_H} \hat{H}_{it} + \varepsilon_{i,t+1}^H \end{aligned}$$

Here, $\varepsilon_{i,t+1}^H$ is a mean-zero i.i.d shock. Resiliency, H_{it}, thus behaves similarly to an AR(1) process that mean-reverts to H_{i*} and does so at a rate that is approximately ϕ_H. The variable part of resiliency, $\hat{H}_{it}$, is modeled separately simply out of convenience. While the functional form of $\hat{H}_{it}$ may seem peculiar as an AR(1) "like" process, the reasons for this modeling choice lie in the mechanics of LG processes. Assuming resiliency itself follows an LG-process ensures prices in the economy are linear in resilience; hence, resilience is the key state variable for determining price-dividend ratios. Now, I will present some results on equilibrium asset prices that will be used in this paper.

4.3 Theoretical Results for Stocks and Options

Result 1 (Gabaix (2012)). *The equity premium (conditional on no disasters) is:*

$$\begin{aligned} r^e_{it} - r_{ft} &= \delta - H_{it} - r_{ft} \\ &= p_t \, \mathbb{E}_t[B^{-\gamma}_{t+1}(1 - F_{i,t+1})] \end{aligned} \tag{2}$$

where r^e_{it} is the expected return of the asset conditional on no disasters and $r_{ft} = \delta - p_t \, \mathbb{E}_t[B^{-\gamma}_{t+1} - 1]$ is the risk-free rate in the economy conditional on no disasters. As expected, assets that are more resilient through crises command a lower equity premium.

Due to the properties of the model, it can be shown that the price of stock i evolves according to:

$$\frac{P_{t+1}}{P_t} = e^{\mu_{it}} \times \begin{cases} e^{\sigma u_{i,t+1} - \sigma^2/2} & \text{if there is no disaster at } t+1 \\ F_{i,t+1} & \text{if there is a disaster at } t+1 \end{cases} \tag{3}$$

where $u_{i,t+1}$ is a standard Gaussian variable and μ_{it} is determined in equilibrium.[12] For short horizons the level of price growth will have trivial effects on the asset prices and I will not spend more time discussing its properties. Intuitively, this specification says that the most of the time we live in log-normal world for stock prices but sometimes there is a disaster and the stock price falls a dramatic amount.

[12] When I extend the model to accommodate Epstein-Zin preferences in Section 4.5, I will provide more detail concerning equilibrium prices and the consistency of return processes in this economy. I skip over it now for brevity.

A Comment on the Price Specification

As a quick check of how valid the assumption of log-normality is outside of disasters, I looked at monthly returns of the CRSP Value Weighted Index from January 1945 to June of 2013. For the entire sample, the unconditional skewness is -0.51; however, if I exclude the ten months with the most negative returns the skewness is only slightly below zero at -0.07. This is comforting from the model's perspective since this indicates there are only a few extreme observations that are driving the skewness of the full sample. Among the ten most negative returns are well-known crises such as October of 2008 and October of 1987. Additionally, the truncated time-series is only slightly leptokurtic with a kurtosis of 3.28. Moreover, it is also well-known that the market index demonstrates more negative skewness than the average stock, so taking these metrics as representative of any generic stock i is on the conservative side. I think it is therefore safe to at least assume that log-returns are roughly symmetric in normal times. Allowing a more general symmetric distribution than log-normality doesn't change the results, but makes the math substantially more complicated, while the intuition of my measure remains the same. Hence, I conclude that this assumption is innocuous for the purposes of my paper.

Options in a Disaster Economy

I also consider the price of a one-period European put option on stock i with *strike to spot ratio* K. Denote the value of this put option as $V_t^{put}(K) = \mathbb{E}_t[\frac{M_{t+1}}{M_t} \cdot \max(0, K - P_{i,t+1}/P_{i,t})]$. The price of a put option in this economy is then easily derived according to the following result.

Result 2 (Gabaix (2012)). *The value of a put with a one-period maturity* V_{it}^{put} *is*

$$\begin{aligned} V_{it}^{put}(K) &= V_{it}^{ND,put}(K) + V_{it}^{D,put}(K) \\ V_{it}^{ND,put} &= e^{-\delta+\mu_{it}}(1-p_t)V_{it}^{BS,put}(Ke^{-\mu_{it}}, \sigma) \\ V_{it}^{D,put} &= e^{-\delta+\mu_{it}} p_t \, \mathbb{E}_t[B_{t+1}^{-\gamma} \cdot \max(0, Ke^{-\mu_{it}} - F_{i,t+1})] \end{aligned} \tag{4}$$

where $V^{BS,put}(K, \sigma)$ *is the Black-Scholes value of a put with strike* K*, volatility* σ*, initial price* 1*, maturity* 1 *and interest rate zero.*[13]

Thus, a put option in this economy is made up of two parts: a disaster component and a normal times component that corresponds to the Black-Scholes price. Each of these components is naturally weighted by their respective probabilities.

4.4 A Measure of Disaster Risk

Now I will derive the value of what I will call disaster risk DR in this economy. The concept of disaster risk is most often found in foreign exchange derivatives, where in this market it is commonly referred to as a risk-reversal. In currency markets, a risk-reversal simultaneously purchases an OTM put option and sells a symmetrical OTM call option on exchange rates. Risk-reversals are often used by traders speculating on currency movement and also used to bound gains and losses in carry trades. This strategy is, by design, exposed to large movements in exchange rates. I will extend the concept of risk-reversals to my DR measure and show that DR_{it} has the special property of canceling out the Black-Scholes component

[13] As a matter of interpretation, the one-period nature of this option is not necessarily a one-year option. We can think of this option as expiring the next time the dividend is paid and over that time period the risk-free rate is zero, and the volatility over this time period is σ. I will be considering one-month options so this seems like a reasonable approximation. An analogous setup is derived for currency options in Farhi, et al (2009).

of both a call and a put, thereby leaving only the portion of their prices that comes from the probability of a rare disaster. Define disaster risk as follows:

Definition 1. *The disaster risk for firm i at moneyness M is defined as:*

$$DR_{it}(M) = Put(M) - Me^{-\mu_{it}} Call(M^{-1} e^{2\mu_{i,t}})$$

In the limit of small time intervals, it can be written more succinctly as:

$$DR_{it}(M) = Put(M) - M \cdot Call(M^{-1}) \tag{5}$$

The following proposition derives the measure of disaster risk in this economy.

Proposition 1. *Assuming $M > F_{i,t+1}$ a.s and $MF_{i,t+1} < 1$ a.s., the level of disaster risk (in the limit of small time intervals) is:*

$$DR_{i,t}(M) = p_t \, \mathbb{E}_t[B_{t+1}^{-\gamma}(M - F_{i,t+1})] \tag{6}$$

Proof. The proof is in Appendix A. □

The two assumptions that deliver Proposition 1 are rather innocuous: $M \overset{a.s}{>} F_{i,t+1}$ states that the dividend process drops by a sufficient amount in disasters (or rather the put option is not too far out of the money) and $MF_{i,t+1} \overset{a.s}{<} 1$ states that the put option leg of DR_{it} is not too far in the money. Hereafter, I will only explicitly write the functional dependence of $DR_{it}(M)$ on the moneyness M when it is needed for clarity. The key feature of this economy that DR_{it} takes advantage of is the fact that with some non-zero probability we live in a Black-Scholes type world. It is precisely this feature that lets the Black-Scholes put and call components of disaster risk cancel each other out, leaving only the components of the options that are featured in non-normal times. In this economy this is exactly the state of the world when there is a consumption disaster. Furthermore, it is clear from Proposition 1 that a portfolio of puts and calls can together capture tail risk - very deep OTM puts are not the only way to do so.[14]

There are two natural corollaries that come from Proposition 1. The first links the equity premium exactly to the sum of $DR_{it}(M)$ across different M and the second links the risk-neutral probability of disaster to the difference in $DR_{it}(M)$ for different M.

4.4.1 The Equity Premium and $DR_{it}(M)$

Corollary 1. *Consider two measures of disaster risk $DR_{it}(M)$ for the same firm at M_1 and M_2 such that $M_1 + M_2 = 2$. Then the equity premium (conditional on no disasters) can be deduced using the price of these two measures as follows:*

$$\begin{aligned} r^e_{it} - r_{ft} &= p_t \, \mathbb{E}_t[B^{-\gamma}_{t+1}(1 - F_{i,t+1})] \\ &= \frac{DR_{it}(M_1) + DR_{it}(M_2)}{2} \end{aligned} \tag{7}$$

Proof. The proof is straightforward from the measure of a disaster risk derived in Proposition 1. □

Notice this amounts to constructing one DR measure with OTM put/calls and one DR measure with ITM puts/calls. The value of both measures still follows Proposition 1 as long as disaster measures are not too far OTM or ITM. I will henceforth call the equity premium derived from disaster risk measures as the "DR implied equity premium". Corollary 1 gives

[14] In fact, the theory suggests that if the put is too deep OTM then $M < F_{i,t+1}$ and the disaster component of the put price goes to zero.

a very easy way to back out the equity premium for stocks from a combination of options on the stock. The construction of my DR measure would hold in any economy in which there is a probability of a Black-Scholes world. However, it is worth noting that the link to the ex-ante equity premium is highly dependent on the model set up and I will only rely on it as a qualitative check of my disaster risk measure.

4.4.2 The Probability of Disasters and $DR_{it}(M)$

Corollary 2. *Consider two DR_{it} measures at M_1 and M_2 such that $M_2 > M_1$. Then the risk-neutral probability of a disaster can be deduced as follows:*

$$p_t \, \mathbb{E}_t[B_{t+1}^{-\gamma}] = \frac{DR_{it}(M_2) - DR_{it}(M_1)}{M_2 - M_1} \tag{8}$$

Proof. The proof is straightforward from the expression for disaster risk derived in Proposition 1. □

Corollary 2 delivers an easy way to use DR_{it} at different moneyness to back out the risk-neutral probability of a disaster.[15] Notice this is not a true probability because I have chosen not to normalize it, but it still contains the relevant information about the physical probability of a disaster and risk-aversion. To be precise, $p_t E_t[B_{t+1}^{-\gamma}]$ is the state-price for the disaster state. When the state-price increases, investors are willing to pay a higher price for a contract that pays \$1 if there is a disaster and nothing otherwise. A combination of DR_{it} for the same firm provides a straightforward way to isolate this state-price. Two disaster risk measures on the same firm will have the same firm-specific component $F_{i,t+1}$ so it makes sense that subtracting the two will filter this portion out and leave a quantity that is related

[15] Unlike the implied equity premium, I do not expect the risk-neutral probability of disasters implied by DR_{it} to be so "model-specific". This follows directly from the unspecified nature of $F_{i,t+1}$.

only to aggregate disaster. Additionally, the probability of a disaster in this case must always be linked to the time to expiration of the option. For example, if the disaster risk measure is constructed using 30-day options, then the probability (and also implied equity premium) of a disaster that is derived from DR_{it} is naturally over a 30-day window as well.

Now that I have shown theoretically that my DR_{it} measure contains valuable information about equity returns and consumption disasters, I turn to connecting the price dynamics of a given firm to that firm's disaster risk. This requires extending the model to accommodate Epstein-Zin preferences.

4.5 Price Dynamics, Disaster Risk, and Epstein-Zin Preferences

In order to build testable predictions from my measure of disaster risk, it is useful to consider how DR_{it} should relate to future prices. Thus far, I have worked out of an economy with power utility since most of the mathematics is simpler in this setting, while retaining most of the intuition of the model. The Epstein-Zin (EZ) extension, however, is necessary to understand the price dynamics in this economy fully. In the power utility case an increase in the probability of disaster can actually cause the price level to *increase* because there is such a strong effect on the risk free rate. The Epstein-Zin extension solves this problem by decoupling risk-aversion from the intertemporal elasticity of substitution. In this section, I first present what equilibrium prices are in an economy with Epstein-Zin preferences. Next, I show how DR_{it} still captures disaster risk in an EZ economy, and in fact remain unchanged from the power utility case. An important implication of this result is that the risk-neutral probability of disaster is inferred from my measure exactly as in Corollary 2. Finally, I link realized returns to changes in the probability of disaster. The Epstein-Zin extension is crucial here in order to ensure that increases in the probability of disaster coincide with decreases in prices, as opposed to increases in prices as is the case with power utility.

I will present a first-order approximation of an economy with EZ preferences.[16] By a first-order approximation, I mean that I use an approximation of the stochastic discount factor. The fundamentals of the economy remain the same as before, with the only change being that the representative investor now has Epstein-Zin preferences. In their seminal paper, Epstein and Zin (1989) show that the stochastic discount factor is given by:

$$\frac{M_{t+1}}{M_t} = e^{-\rho/\chi} \frac{C_{t+1}}{C_t}^{-1/(\chi\psi)} R_{c,t+1}^{1/\chi - 1} \tag{9}$$

where $R_{c,t+1} = P_{c,t+1}/(P_{ct} - C_t)$ is a return to a claim on aggregate consumption. The decoupling index of Epstein-Zin preferences is $\chi \equiv (1 - 1/\psi)/(1 - \gamma)$ where ψ is the IES and γ is risk-aversion. When pricing a claim to consumption, it is natural to define the resilience of this claim. The power utility resilience of the consumption claim is then analogously defined as:

$$H_{C,t} = p_t \, \mathbb{E}_t[B_{t+1}^{1-\gamma} - 1]$$

In his online technical appendix, Gabaix (2012) shows how the following approximation holds to a first-order:

$$\begin{aligned} \frac{M_{t+1}}{M_t} &= e^{-\rho/\chi} \frac{C_{t+1}}{C_t}^{-1/(\chi\psi)} R_{c,t+1}^{1/\chi - 1} \\ &\approx e^{-\delta}(1 + (\chi - 1) H_{C,t} + \epsilon_{t+1}^M) \times \begin{cases} 1 & \text{if there is no disaster at } t+1 \\ B_{t+1}^{-\gamma} & \text{if there is a disaster at } t+1 \end{cases} \end{aligned}$$

[16] The fully solved model with no approximations is easiest to do in continuous time, which is why I rely on the first order approximation here.

Writing the stochastic discount factor as a linear function of consumption resilience $H_{C,t}$ will prove useful when trying to price the option components of the disaster risk measure.

In the power utility case, the key state variable for price-dividend ratios was the resilience of the stock H_{it}. Analogously, the central ingredient for determining prices in this economy is now Epstein-Zin enriched resilience, which is defined as $H_{it}^{EZ} = H_{it} + (\chi - 1)H_{C,t} = p_t \, \mathbb{E}_t[B_{t+1}^{-\gamma}(F_{i,t+1} + (\chi - 1)B_{t+1}) - \chi]$ for firm i. In the power utility case, $\chi = 1$ and EZ-enriched resilience collapses to the original definition of resilience. The additional term $(\chi - 1)H_{Ct}$ comes from addition of the same term to the stochastic discount factor. I can rely on the same LG-process methodology as before to solve for the value of a claim to the dividend stream for a given firm i.[17]

Result 3 (Gabaix (2012))**.** *The stock price in the Epstein-Zin case is given by:*

$$P_{it} = \frac{D_{it}}{\delta_i}\left(1 + \frac{\hat{H}_t^{EZ}}{\delta_i + \phi_H}\right) \tag{10}$$

where $\hat{H}_{it}^{EZ} = H_{it}^{EZ} - H_{i*}^{EZ}$ *and* $H_{i*}^{EZ} = p[\bar{B}^{-\gamma}(\bar{F}_i + (\chi - 1)\bar{B}) - \chi]$ *is the mean level of EZ-resilience to which resilience reverts to.*

Based on the definition of resilience it is easy to see that, as in the power utility economy, assets with higher resilience fare better in crises than assets with low resilience. The next natural step is to derive my measure of disaster risk in the economy with EZ preferences. It turns out that DR_{it} remain largely unchanged.

Proposition 2. *Assuming* $M > F_{i,t+1}$ *a.s and* $MF_{i,t+1} < 1$ *a.s., the level of disaster risk*

[17] In the appendix, I demonstrate the general intuition for using LG-processes, which can be applied to both the power utility case or the EZ-case.

(in the limit of small time intervals) in an economy with EZ-preferences is:

$$DR_{i,t}(M) = p_t \, \mathbb{E}_t[\eta_t B_{t+1}^{-\gamma}(M - F_{i,t+1})] \tag{11}$$

where $\eta_t \equiv (1 + (\chi - 1)H_{ct})$. Here, H_{ct} is the power-utility resilience of a "stock" that pays aggregate consumption as its dividend.

Proof. The proof is in Appendix A. □

The intuition from the power utility case still applies in an economy with EZ preferences - DR_{it} cancels out the "normal-times" risk and reflects only the components of the economy when there is a disaster. The only difference between DR_{it} in a power utility world and an EZ world is the factor η_t, which is simply a correction for the EZ-preferences. In the appendix, I show that η_t is very close to 1 and thus can be ignored:

$$DR_{i,t}(M) \approx p_t \, \mathbb{E}_t[B_{t+1}^{-\gamma}(M - F_{i,t+1})] \tag{12}$$

Importantly, Corollary 2 still holds and enables me to use DR_{it} to back out the risk-neutral probability of disaster.[18]

I've shown how the Epstein-Zin extension of the disaster model has virtually no effect on the disaster content of my measure, and have also presented equilibrium stock prices in this setting. Now I turn to developing a link between realized returns and the level of disaster risk, since this is one way I will empirically test the model. The EZ-model says

[18] In the EZ world, the risk-neutral probability of disaster is formally $p_t \, \mathbb{E}_t^D[\eta_t B_{t+1}^{-\gamma}]$. Since $\eta_t \approx 1$, I refer to $p_t \, \mathbb{E}_t^D[B_{t+1}^{-\gamma}]$. Moreover, η_t is very slow moving and thus time-variation in the risk-neutral probability of disaster will not be affected by ignoring it.

that an increase in the probability of disaster are accompanied by a decrease in prices, with the amount of price sensitivity being related to a stock's disaster risk. To see why, iterate Equation (10) one period forward and rearrange terms:[19]

$$\begin{aligned}\log\left(\frac{P_{i,t+1}}{P_{it}}\right) &= \log\left(\frac{D_{i,t+1}}{D_{it}}\right) + \log\left(1+\frac{\hat{H}^{EZ}_{i,t+1}}{\delta_i+\phi_H}\right) - \log\left(1+\frac{\hat{H}^{EZ}_{it}}{\delta_i+\phi_H}\right) \\ &\approx \log\left(\frac{D_{i,t+1}}{D_{it}}\right) + \left(\frac{1}{\delta_i+\phi_H}\right)\left(\hat{H}^{EZ}_{i,t+1} - \hat{H}^{EZ}_{it}\right) \end{aligned} \tag{13}$$

Over a small time limit, we can think of variation in resilience (and therefore price-dividend ratios) as coming from variation in the probability of disaster. Hence, I will treat $\mathbb{E}^D_t[B^{-\gamma}_{t+1}(F_{i,t+1}+(\chi-1)B_{t+1})-\chi]$ as constant over short periods of time. After substituting the definition of EZ-enriched resilience $\hat{H}^{EZ}_{i,t+1}$ into Equation (13), the realized log-return of the stock can then be written as:

$$\begin{aligned}\log\left(\frac{P_{i,t+1}}{P_{it}}\right) &= \log\left(\frac{D_{i,t+1}}{D_{it}}\right) + \left(\frac{1}{\delta_i+\phi_H}\right)(p_{t+1}-p_t)\underbrace{\mathbb{E}^D_t[B^{-\gamma}_{t+1}(F_{i,t+1}+(\chi-1)B_{t+1})-\chi]}_{\equiv\Gamma_i} \\ &= \log\left(\frac{D_{i,t+1}}{D_{it}}\right) + \left(\frac{\Gamma_i}{(\delta_i+\phi_H)(\mathbb{E}_t[B^{-\gamma}_{t+1}]}\right)(p^*_{t+1}-p^*_t) \\ p^*_{t+1} &\equiv p_{t+1}\,\mathbb{E}_t[B^{-\gamma}_{t+1}] \end{aligned} \tag{14}$$

Here p^*_{t+1} is simply the risk-neutral probability of a disaster, which I can empirically estimate. Equation (14) is a regression of realized returns (or realized changes in the price-dividend ratio) on changes in the risk-neutral probability of disaster. In order to make sure the regression coefficient on $(p^*_{t+1}-p^*_t)$ is constant, I assume B_{t+1} is stationary and Γ_i is a

[19] The reader may find the statement in Equation (3) contradictory to the following derivation. However, recall that the flexibility in the noise process for dividends and resilience makes Equation (3) and (13) consistent with each other. In turn, the model implied representation of DR_{it} still holds. I show this formally in the proof of Proposition 2 in Appendix A

constant. The assumption that B_{t+1} is stationary ensures that the expected severity of consumption disasters does not change through time. Furthermore, a constant expected disaster severity implies variation in the state price of a disaster are driven by variations in p_t, which is a harmless assumption. Moreover, by assuming Γ_i to be a firm specific constant, I am simply positing that the covariance between marginal utility and $F_{i,t+1}$ is constant through time.[20]

According to the calibration in Gabaix (2012) and Barro (2006), $\chi < 0$, which is necessary in order for the stock price to *drop* as the probability of disaster *increases*. In the power utility case, the regression coefficient on $(p^*_{t+1} - p^*_t)$ in (14) would lead to the prediction that *increases* in the probability of a disaster coincides with an *increase* in prices and a positive realized return – hence the need for the Epstein-Zin extension. So, under reasonable calibrations $\Gamma_i < 0$ and the disaster model implies that a regression of realized returns on changes in p^*_{t+1} should yield a regression coefficient that is negative. What about the magnitude of this regression? There are two firm specific constants δ_i and Γ_i, as well as a economy wide parameters ϕ_H and $E_t[B^{-\gamma}_{t+1}]$; hence, cross-sectional variation in the regression coefficient will be driven by cross-sectional variation in Γ_i and δ_i. Here, δ_i is the effective discount rate of the stock in good times, which will be relatively homogeneous across firms compared to Γ_i, since Γ_i reflects the recovery rate F_i of the firm in a disaster. It is straightforward to see that firms with high resilience will have *less negative* values of Γ_i. Equivalently, firms with high DR_i will have *more negative* Γ_i. The preceding logic then summarizes nicely in the following empirical prediction:

Empirical Prediction 1. *In a regression of realized-returns on changes in the risk-neutral*

[20] In fact, Gabaix (2012) argues that most of the time variation in resilience comes from p_t, and that we can treat $F_{i,t+1}$ as a constant. Assuming B_{t+1} to be strongly stationary does not conflict with resilience as an LG-process, since there is still modeling flexibility in p_t. This simplification will prove useful when evaluating the empirical implications of this model.

probability of disaster, the coefficient should be significant and negative. Moreover, for stocks with higher resilience (lower DR_i), the coefficient on the regression should be closer to zero (i.e. stocks with low disaster risk are less sensitive to changes in the probability of disaster).

Here I am assuming most of the variation in realized-returns is driven by changes in the probability of disaster and not by growth rate of dividends. An alternative prediction which I will test in Section 6 is that changes in the price dividend ratio should be related to changes in the risk-neutral probability of disasters. In a broad sense, my DR_i measure can be interpreted as measuring the sensitivity of a stock to changes in the probability of a consumption disaster. For a given increase in the probability of a disaster, stocks with high disaster risk will experience very sharp decreases in prices compared to stocks with low disaster risk. This is because stocks with high DR_i have more negative Γ_i's in the price sensitivity regression (14).

An Alternative Interpretation of DR_{it}

Alternatively, DR_{it} can be thought of as a measuring covariance of a stock with consumption during a disaster. To see why, consider two firms i and j whose disaster risk measures at time t are given by:

$$\begin{aligned}
DR_{it}(m)/p_t &= \mathbb{E}_t[B_{t+1}^{-\gamma}(M - F_{i,t+1})] \\
&= (m - E_t[F_{i,t+1}])E_t[B_{t+1}^{-\gamma}] - cov_t(B_{t+1}^{-\gamma}, F_{i,t+1}) \\
DR_{jt}(m)/p_t &= \mathbb{E}_t[B_{t+1}^{-\gamma}(M - F_{j,t+1})] \\
&= (m - E_t[F_{j,t+1}])E_t[B_{t+1}^{-\gamma}] - cov_t(B_{t+1}^{-\gamma}, F_{j,t+1})
\end{aligned}$$

Now it is straightforward to see that if stock i's disaster performance has higher covariance with marginal utility in a disaster (i.e. it pays off when the representative agent needs it the most) it will have a *lower* disaster-risk measure. Consider a simplified case where two stocks' disaster performance is on average the same so that $\mathbb{E}_t[F_{i,t+1}] = \mathbb{E}_t[F_{j,t+1}]$. Then in this case, $DR_{it}(m) > DR_{jt}(m) \Leftrightarrow cov_t(B_{t+1}^{-\gamma}, F_{i,t+1}) < cov_t(B_{t+1}^{-\gamma}, F_{j,t+1}) \Leftrightarrow cov_t(B_{t+1}, F_{i,t+1}) > cov_t(B_{t+1}, F_{j,t+1})$ implies that stock i has a higher covariance with consumption during disasters than stock j.[21]

It is also straightforward to show that $DR_{it}(m) > DR_{jt}(m) \Rightarrow \Gamma_i < \Gamma_j$ which maps the intuition of DR_{it} as a measure of consumption disaster covariance to Empirical Prediction 1 - higher disaster risk for firm i over firm j means greater price sensitivity of firm i versus firm j to changes in the aggregate probability of a disaster.[22] To summarize, DR_{it} *captures covariance of a firm with consumption in disasters* and all else equal, firms with high DR_{it} have high covariance with the consumption disaster process B_{t+1}. Importantly, historical data will provide poor estimates of a firm's covariance with consumption during disasters since these episodes are so rare. Even if reliable estimates were available, firm characteristics change over time making options an even more attractive place to look for information on future disaster performance.

I also want to emphasize how my simple disaster risk measure allows me to capture a disaster covariance in a way that an simple OTM put could not. In the proof of Proposition 2, it is easy to see that an OTM put will contain a term capturing the covariance of B_{t+1} an $F_{i,t+1}$, but will also contain a term capturing risk in normal times. DR_{it} then provides the only way to isolate the term containing disaster risk. In fact, if I were to augment a non-disaster stochastic process to consumption growth then the OTM put would be even

[21] To see why, in a disaster $B_{t+1} < 1$ and $\gamma > 1$.

[22] In fact, as long as $F_i < F_j < 1$ (i.e. both stocks drop in a disaster), $\Gamma_i < \Gamma_j < 0$.

"noisier" since it will contain a term capturing covariance with the normal times risk-factor. DR_{it}, however, is able to focus more on disaster risk.

5 Data and Methodology

The primary source of data for this paper is the Options Metrics Volatility Surface (OMVS) available from the WRDS database. I use thirty-day European options on constituents of the S&P 500 Index. In the previous sections, I was careful to use only members of the index that were present at the time of portfolio formation in order to avoid a lookahead bias. For what remains, however, I relax this constraint since I will want to use the most liquid options. The short-dated options are used because much of the theory developed in Section 4 relies on short time to maturity arguments. In addition, these options are less likely to suffer from liquidity biases and therefore produce spurious results. The sample of options runs from January 1996 to January 2012. The OMVS is formed daily and contains options of varying fixed maturities and implied strikes. The volatility surface is created using a kernel smoothing technique. A full description of the OMVS can be found on the Options Metrics website. For a detailed explanation of how all variables (my disaster risk measure, implied equity premium from DR_{it}, and the risk-neutral probability of disaster) are constructed please refer to Appendix B.2.

6 Empirical Tests of the Disaster Model

The disaster model in Section 4 delivered two broad implications:

1. There is a common aggregate disaster probability p_t that affects all stocks, including

the market. This (risk-neutral) probability can be inferred using the DR measure of *any* stock according to Corollary 2.

2. Relative to firms with low disaster risk, firms with higher DR measures should exhibit a stronger *negative* contemporaneous return relationship with changes in the probability of a disaster. This is because when the probability of a disaster increases, firms with high DR should experience large price drops.

Both of these implications are testable empirically, which I now treat in turn.

6.1 The Risk-Neutral Probability of a Disaster

Corollary 2 says that the risk-neutral probability of a disaster can be inferred using the DR measures of a given firm at each point in time. Thus, I construct a daily time-series of disaster probabilities at each point in time, inferred from each firm. From here, I aggregate daily measurements for each firm into monthly measurements by simply taking the within-month median of the daily measures. This procedure delivers me a monthly disaster probability p^*_{it} from each firm i. Due to measurement noise in the cross-section, I then posit the following factor structure for the time-series of p^*_{it}:

$$p^*_{it} = p^*_t + e_{it}, \qquad i = 1, ..., N$$

where each e_{it} is i.i.d noise. I then use standard latent factor analysis to estimate the common factor p^*_t, which here represents the aggregate risk-neutral probability of a disaster. Bai and Ng (2008) provides an extremely detailed summary of the econometrics of factor analysis methods, including the asymptotic distribution of estimated factors and loadings. I will avoid the distribution theory for factor analysis and instead focus on simpler measures of fit, such as the percent variation captured by the principle components of p^*_{it}. Table 4 summarizes my findings.

Table 4: PCA of the Cross-Section of p^*_{it}

Principle Component	**Percent of Total Variation Captured**
1	66.0%
2	2.1%
3	1.5%
4	1.3%
5	1.2%

Notes: The table reports the results of principle component analysis of the inferred p^*_{it} from a large panel of firms.

Clearly, there is one large principle component that guides much of the variation in p_{it}^* through time, which is exactly what the model suggests. Next, I plot the extracted factor based on the preceding principle component analysis:[23]

Figure 2: Aggregate Disaster Probability (up to a constant scale)

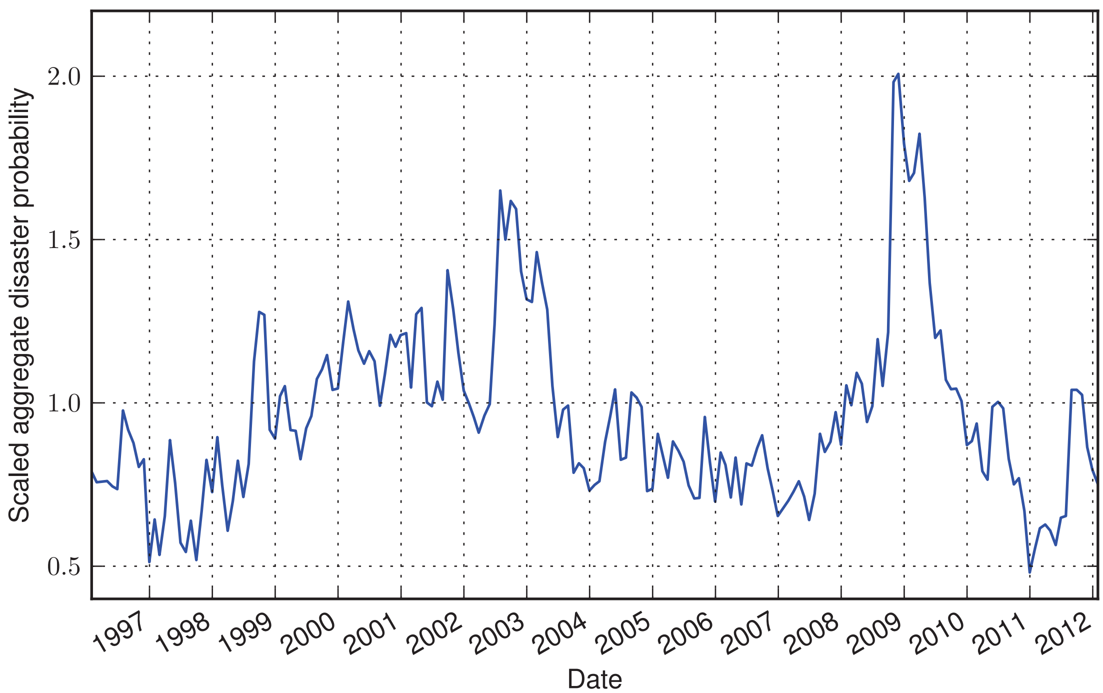

Notes: This figure plots the factor extracted from the panel of p_{it}^* from January 1996 to January 2012. The factor is interpreted as the probability of disaster common to all firms, and is estimated up to a scaling constant.

[23]Note in Bai and Ng (2008), there is substantial time dedicated to choosing the right number of latent factors k. Since there is such as large principle component in this case, their selection method will deliver a single factor (i.e $k = 1$) as it is based on the sum of squared error terms e_{it}.

The scale on the y-axis of Figure 2 does not correspond to a probability. This is because when conducting factor analysis, it is only possible to determine the factors and loadings up to a linear rotation. Practically, this means that I can pin down the "clean" aggregate probability of a disaster up to a scaling constant; however, I will shortly explore an alternative method so I can say something about the absolute likelihood of a disaster. Henceforth, I will refer to the common factor pulled out of the individual disaster probabilities p^*_{it} as the aggregate disaster probability and denote it by p^*_t. The time-series of disaster probability evolves as one might expect - the spikes correspond to some well-known "crises" over the sample period. The spike in November 1997 corresponds to the Asian currency crisis and its spread to Indonesia and South Korea, and in September 1998 we observe the LTCM crisis. The terrorist attack on the United States in September 2001 also resulted in a jump in disaster probabilities. In July 2002, WorldCom filed for bankruptcy which at the time was the largest corporate insolvency ever, there was increased fighting on the Gaza Strip, and an earthquake occurred in Germany; all of these events likely contributed to the jump in p^*_t at the end of this month. The probability of disaster also spiked later that year in September 2002 as the US announced its plan to invade Iraq. The low volatility period from 2002 to 2007 corresponds to a relatively low value for p^*_t, with disaster risk rising at the onset of the Great Recession 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. There was another spike in May 2010 following the 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 occurred in August 2011, a month which saw the concerns about the sovereign debt crisis worsen and a downgrade of the creditworthiness of U.S. debt.

While p_t^* captures the time-series trends in disaster probability, we might also be interested in the actual magnitude of the probability of disaster. As such, I alternatively use DR measures on the S&P 500 market portfolio to infer what the market itself suggests as the probability of a disaster. Figure 3 plots the disaster probability inferred from the S&P 500, denoted $p_{spx,t}^*$ and a scaled version of p_t^* such that the means of the two series are equal. Moreover, I have annualized the monthly measures of disaster probability simply by multiplying by 12. The first thing to notice is the magnitude of $p_{spx,t}^*$. The unconditional average of $p_{spx,t}^*$ is 21.8%, which is very close to the calibration in Gabaix (2012) of 19.2%. Granted, Gabaix's (2012) calibration are designed to match the true p_t^*, but insofar as the market is a good indication of the aggregate economy, we can use $p_{spx,t}^*$ without worry. To this end, the correlation between p_t^* and $p_{spx,t}^*$ is quite high at 70%, which is further confirmation that p_t^* is working at a macro level and the two series are capturing roughly the same thing. In the Online Appendix, I also validate further that p_t^* operates at an aggregate level by demonstrating its forecasting power for macroeconomic aggregates such as changes in unemployment and investment activity.

Unsurprisingly, the noise in p_t^* is much lower than $p_{spx,t}^*$. The reason is that p_t^* is estimated in a way that minimizes the influence of the measurement noise $e_{i,t}$, whereas $p_{spx,t}^*$ still contains its own measurement noise. In fact, the ratio of the mean to the standard deviation for p_t^* is 3.45 compared to that of 1.91 for $p_{spx,t}^*$. Therefore, to simultaneously take advantage of the precision of p_t^* and the magnitudes of $p_{spx,t}^*$, I will use the scaled version of p_t^* whose unconditional average matches that of $p_{spx,t}^*$ for the remainder of the paper.

Figure 3: Aggregate Disaster Probability (scaled)

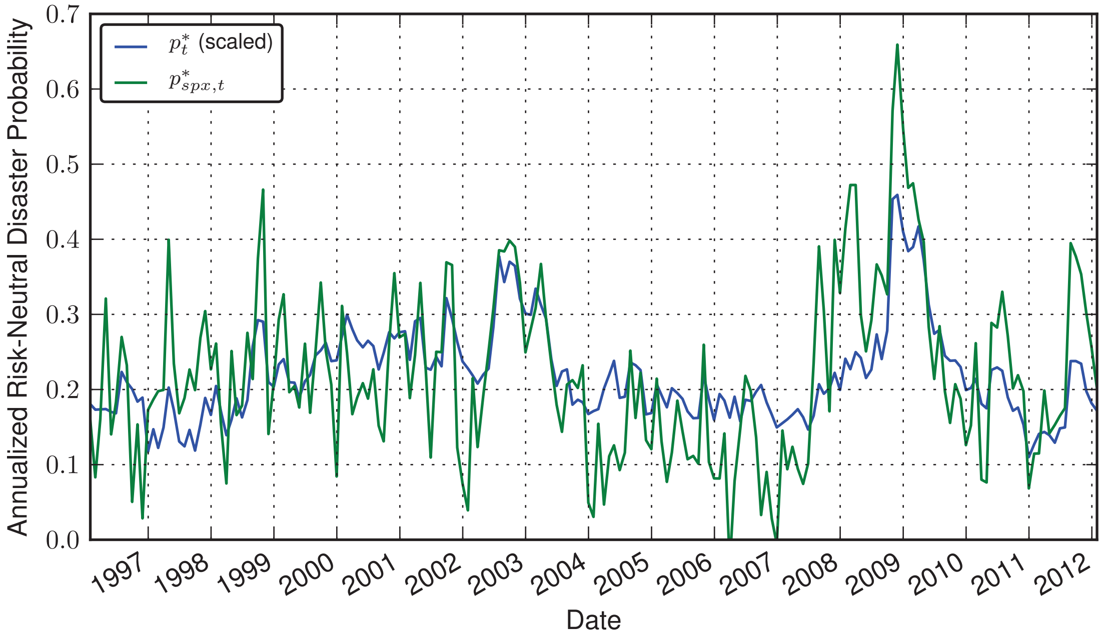

Notes: This figure plots the common factor p_t^* extracted from the panel of p_{it}^*, as well as the probability of disaster implied by S&P 500 DR_{spx}. The factor from the cross-section of p_{it}^* is normalized so its average matches that of $p_{spx,t}^*$. The sample is from January 1996 to January 2012

6.2 Realized Returns and Changes in the Probability of Disaster

The discussion in Section 4.5 suggested a tight link between realized returns and changes in the risk-neutral probability of disasters. Since the probability of disaster is backed out from combinations of my DR_{it} measure, these empirical predictions also give credence to the notion that DR picks up on the important asset pricing aspects of disasters in the economy. I will first explore the link between returns and disasters using the S&P 500, and then I will further verify this connection further by looking at heterogeneity of DR_i in the cross-section and realized returns.

6.2.1 Realized S&P 500 Returns and the Probability of Disaster

Empirical Prediction 1 suggests a simple regression of realized-returns on changes in the risk-neutral probability of disaster should yield a statistically significant and negative coefficient. While this holds true of any firm or portfolio, I start by testing this for the S&P 500 index.

Table 5: Realized-Return versus Probability of Disaster Regressions

Panel A - Realized One-Month S&P 500 Returns

Constant	0.003
	(1.22)
Δ(Risk-Neutral Probability of Disaster)	-8.13**
	(-6.14)
R^2	20.3%

Panel B - Change in log(Price-Dividend Ratio) for S&P 500

Constant	7e-4
	(0.24)
Δ(Risk-Neutral Probability of Disaster)	-8.58**
	(-6.49)
R^2	22.0%

** Denotes significance at a 5% confidence interval

Notes: The table reports results from regressions of S&P 500 returns over a one month horizon from the period of January 1996 to January 2012. All standard errors are calculated according to Newey-West (1987). T-statistics for each regression are listed in parenthesis. In Panel A, the LHS variable is one-month realized log-returns of the market. In Panel B, the LHS variable is the change in log(price-dividend) ratio.

Panel A of Table 5 regresses realized market returns on changes in the aggregate probability of a disaster. The coefficient in this regression is significantly negative with a Newey-West t-statistic of -6.14. Following with intuition, increases in the probability of disaster should be accompanied by a decrease in prices since the required premium to hold equity increases - the price must drop to compensate investors for the increased disaster probability. Panel B repeats the regression, but uses changes in the log(price-dividend ratio) as the LHS variable. The regression results remain nearly identical, which suggests that the variation in the dividend payout are not driving the changes in the price.

To access the economic significance of the regression, let us return to Equation (14), which I restate below:

$$\log\left(\frac{P_{m,t+1}}{P_{mt}}\right) = \log\left(\frac{D_{m,t+1}}{D_{mt}}\right) + \left(\frac{\Gamma_m}{(\delta_m + \phi_H)(\mathbb{E}_t[B_{t+1}^{-\gamma}]}\right)(p_{t+1}^* - p_t^*)$$

$$\Gamma_m = \mathbb{E}_t^D[B_{t+1}^{-\gamma}(F_{m,t+1} + (\chi - 1)B_{t+1}) - \chi]$$

$$\approx [\bar{B}^{-\gamma}(\bar{F}_m + (\chi - 1)\bar{B}) - \chi]$$

The subscript m emphasizes that I am examining the market. I will follow the calibration in Gabaix (2012) in order to get a feel for the magnitude of the regression coefficient. Specifically, I calibrate the model as follows:

Table 6: Variables Used in Calibration

Variable Name	**Parameter**	**Calibrated Value**
Consumption Recovery	$\bar{B}$	0.66
Market Dividend Recovery	$\bar{F}_m$	0.6
IES	ψ	2
Risk-Aversion	γ	4
EZ Decoupling Index	χ	-1/6
Stocks: Effective Discount Rate	δ_m	0.073/12
Resilience Mean Reversion	ϕ_H	0.13/12

Table 6 says that in the event of a disaster aggregate consumption drops by 34% and the aggregate stock market index drops by 40%, over a *one year* horizon. Thus, at a monthly frequency we see a 2.8% drop in consumption and a 3.3% drop in the stock market. The reason I keep these variables in annualized terms is because the estimated p^*_{t+1} is backed out using monthly options and is in essence divided by 12 already.[24] Since I set Γ_m to be a constant, I use the unconditional values of all stochastic variables in its definition to get a rough idea of magnitude. An equally appealing interpretation of this calibration would be to set $\mathbb{E}_t[B^{-\gamma}_{t+1}F_{i,t+1}]$ to its own constant. Simple calculation shows that the coefficient on changes in the risk-neutral probability of disaster should be:

$$\left(\frac{\Gamma_m}{(\delta_m + \phi_H)(B^{-\gamma}_{t+1})}\right) \approx -8.17$$

Under reasonable calibration parameters for the model, the model implied regression coefficient is remarkably close to the empirical regression coefficient of -8.13.

6.2.2 Realized Firm Returns and the Cross-Section of DR_{it}

Cross-sectionally, assets with the highest disaster risk should be the most sensitive to changes in the probability of risk-neutral disaster. In order to explore this empirically, I use the time series of DR_{it} described in Section 3 for each of the constituents of the S&P 500 Index as of December 1, 2012. Figure 4 plots the firm-level DR data, aggregated by industry. For each month and each industry, I simply took the within-month median level of DR across all firms in a given industry. Qualitatively, the results are reasonable as companies in traditionally riskier industries like technology and financial services have higher DR than

[24]This is a heuristic argument. In the continuous time version of the model, the intensity of the jump process is B and the probability of a jump over a short time span is $p\Delta t$.

firms in industries such as industrial goods. The other thing to notice is that the ordering of DR across industries varies through time. For example, at the end of the Internet bubble of the early 2000s, financial firms did not have a large DR, but in the most recent financial crisis this was obviously not true. Clearly the nature of disasters evolves through time, which makes using historical data even less appealing as compared to options which are forward looking.[25]

Figure 4: DR by Industry

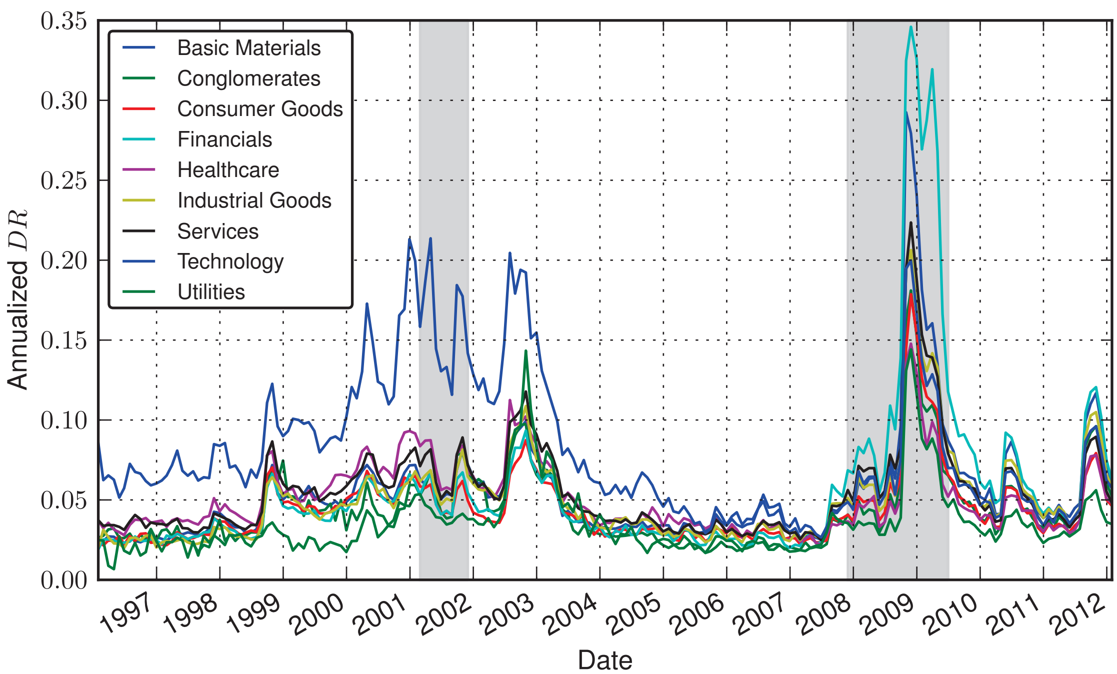

Notes: This figure plots within-month median DR for various industries from January 1996 to January 2012. The grey shading indicates NBER recession dates.

Next, I sort firms based on their median level of DR_{it} over their respective sample period

[25]Indeed, the pattern of DR resembles that of aggregate volatility. However, this is not by construction, since the options used to construct DR are adjusted, at the firm level, to account for differences in volatility across firms. Moreover, as seen in Section 3, sorting stocks based on DR shows almost zero exposure to changes in the VIX. The pattern observed in DR therefore reinforces the intuitive notion that volatility and disaster risk likely influence each other.

and group them into buckets.[26]. The lowest bucket is the set of firms with the lowest median level of DR_{it}, and the highest bucket is the set of firms with the highest median level of DR_{it}. For each firm i in each bucket, I then conduct the following regression:

$$\log\left(\frac{P_{i,t+1}}{P_{it}}\right) = \alpha_i + \beta_i(p^*_{t+1} - p^*_t) \tag{15}$$

where p^*_{t+1} is the aggregate risk-neutral probability of disaster. For each bucket j, I then compute the average regression coefficient across firms in that decile. Formally, I compute $\bar{\beta}_j = \frac{1}{N_j}\sum_{i\in j}\beta_i$, where N_j is the number of firms in bucket j. Empirical Prediction 1 suggests firms with higher DR_{it} should experience larger price drops when p^*_{t+1} increases; in other words, $\bar{\beta}_j$ should be decreasing in j.

Figure 5 summarizes the results of these decile regressions. For each bucket, I plot $\bar{\beta}_j$ as well as the average Newey-West (1987) t-statistic for that bucket. The results confirm the idea that higher tail risk firms are more sensitive to changes in the risk-neutral probability of disaster. The average regression coefficient shows a clear declining pattern when going from the lowest to highest bucket. Moreover, the regression coefficients are for the most part significant at a 10% level only for the higher DR_{it} buckets. Since the p^*_{t+1} used in the regression was calculated using the factor analysis of Section 6.1, it does not necessarily correspond to what the probability of disaster is at the end of each month. Thus, I am not surprised that the regression relationship is not significant for low-disaster risk firms. In summary, the cross-sectional results together with the market-level results demonstrate strong support for the model and the disaster information inherent in DR_{it}.

[26]Hence, the median for each firm is computed using a time series with a daily frequency.

Figure 5: Cross-Sectional Price Sensitivity to Changes in p^*_{t+1}

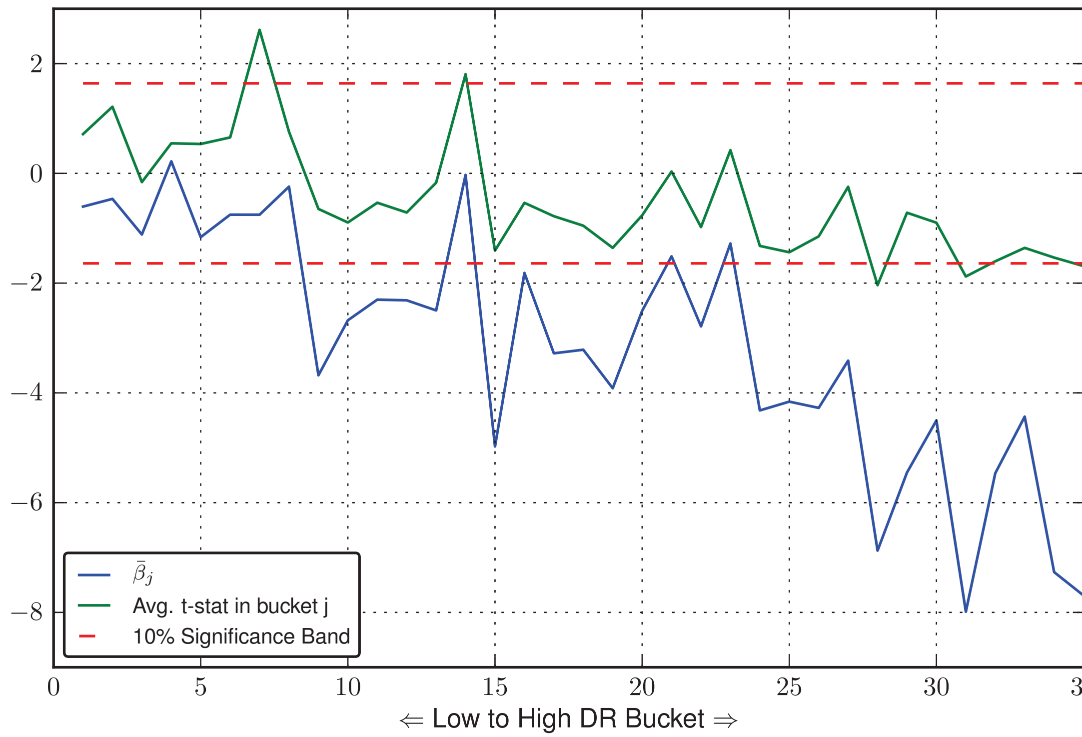

Notes: This figure plots the average regression coefficient of realized returns for a firm against the change in the risk-neutral probability of disaster. Firms have been grouped into buckets based on their median level of DR_{it}. All t-stats were calculated according to Newey-West (1987)

7 Discussion and Conclusion

My Disaster Risk Measure and Measures of Skewness

The early sections of this paper demonstrated that forming portfolios based on DR_{it} delivers large and significant excess returns, even after controlling for standard risk-factors. My disaster risk measure is closely related to the ex-ante risk-neutral skewness of a firm, and indeed the relationship between ex-ante risk-neutral skewness and equity risk premiums has previously explored by other authors. Xing, Zhang, and Zhao (2010) construct a skew measure from options by taking the difference between the implied volatility of an OTM put and an ATM call. They find that forming stock portfolios based on their skew measure sorts generates 10.90% in annualized excess returns after adjusting for Fama-French factors. Another example is Conrad, Dittmar, and Ghysels (2012) who find that stocks that have low, negative ex-ante skewness (implied by options) experience subsequently higher returns than stocks with higher ex-ante skewness. Specifically, they find an excess annualized alpha of 12.12% of a portfolio that is long (short) stocks with low (high) skewness.[27] One key difference in their approach is that they use skewness implied by 3-month options, whereas I form DR_{it} based on 1-month options. Additionally, my measure is constructed using *symmetrical* call and put options, which is crucial according to the theory I develop. On the other hand, risk-neutral skewness in Conrad, Dittmar, and Ghysels (2012) is constructed using the entire cross-section of OTM options and the skewness measure in Xing, Zhang, and Zhao (2010) does not use symmetrical options. Naturally, there will be some overlap between my measure and any measure of skewness; however, I also vary the moneyness of the DR_{it} according to the implied volatility of the firm so my measure will capture third-moments that are in excess of volatility. Furthermore, the economic underpinning for firm-level skewness mattering for equity returns is not totally clear as many of the attempts to explain firm-level skewness and the cross-section of equity returns have relied on behavioral arguments or differences in information between options markets and equity markets.[28] In contrast, my economic model nests into the other asset pricing puzzles addressed in Gabaix (2012) and allows the interpretation of DR_{it} as measuring disaster consumption covariance risk. The

[27] See Table IV, Panel A of their results. Their sample period is also shorter than mine and ranges from 1996 to 2005. They additionally form equity portfolios on a number of different sorting criteria, like double sorts on the intersection of option implied volatility and option implied skewness, and their results are quite robust with regard to skewness.

[28] See for example Barberis and Huang (2008) and Brunnermeier, Gollier, and Parker (2007). Xing, Zhang, and Zhao (2010) put forth an informational argument.

price sensitivity analysis of Section 6 provides additional validation of this interpretation. My model also delivers a measure of the probability of a consumption disaster, whose magnitude is in accordance with much of the disaster literature. Thus, I view the results in the literature regarding firm-level option implied skewness as complimentary to my contribution in this paper.

Concluding Remarks

I have demonstrated that rare consumption disasters indeed imbue themselves into equity prices - a portfolio that is exposed to firms with high disaster risk earns 12.13% in risk-adjusted excess returns. I motivate my measure of disaster risk in a model economy that experiences rare disasters. In this economy, the only source of uncertainty in the stochastic discount factor comes from what happens to aggregate consumption during a disaster; however, the results of the model are not tied to this in any way. Adding a stochastic shock to aggregate consumption in normal times will simply add another term to the risk-premium earned by stocks without changing the key mechanism of the model - that disaster risk is priced into all assets. The determinants of the equity premium in this more general setting would resemble a model with a standard market factor plus an additional factor that captures how the economy performs in a disaster. Instead, the stylized model in this paper focuses on showing theoretically how a DR_{it} can be used to isolate how stocks covary with aggregate consumption during a disaster, thereby providing a measure of disaster risk. In this economy, options are indeed economically redundant assets, yet they provide the econometrician with a richer information set regarding expectations of firm performance during a disaster. Since disasters rarely materialize, measuring firm covariance with disasters based on historical equity returns is extremely difficult. My disaster risk measure is successful in circumventing this problem. Stocks with high DR_{it} are not only more sensitive to changes in the probability of a consumption disaster, but they also earn higher returns on average than stocks with low DR_{it} because of their exposure to disaster risk. Empirically, I find strong support for the model and the disaster content of my measure. Moreover, I find compelling evidence that disaster-risk is priced into the cross-section of equities and therefore DR_{it} can be used to proxy for the measuring disaster risk for a given firm.

This paper focuses on equity derivatives to back out disaster information, but because disasters affect the total assets of a firm there should be analogous information in credit derivatives for a firm. Furthermore, the current model embeds a fixed capital structure for the firm; however, it would be interesting to consider if/how firm decisions depend on the possibility of rare disasters. Additionally, this paper focuses solely on the equity pricing implications of aggregate disaster risk. The model presented in Section 4 has rich implications for bonds as well. In fact, one might think that the yield curve for treasuries would be sensitive only to disasters. In reality, dividend growth of stocks is likely due to many factors so in this sense treasuries may even be a more precise laboratory to study disasters. The variable rare disaster framework provides a tractable way to explore these questions both empirically and theoretically, and is an interesting endeavor for future research.

References

[1] Ang, A., and Bekaert, G. Stock return predictability: Is it there? *Review of Financial Studies 20*, 3 (May 2007), 651–707.

[2] Ang, A., Hodrick, R. J., Xing, Y., and Zhang, X. The cross-section of volatility and expected returns. *The Journal of Finance 61*, 1 (2006), 259–299.

[3] Backus, D., Chernov, M., and Martin, I. Disasters implied by equity index options. CEPR Discussion Papers 7416, C.E.P.R. Discussion Papers, Aug. 2009.

[4] Bansal, R., and Yaron, A. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance 59*, 4 (2004), 1481–1509.

[5] Barberis, N., and Huang, M. Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review 98*, 5 (September 2008), 2066–2100.

[6] Barro, R. J. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics 121*, 3 (2006), 823–866.

[7] Black, F., and Scholes, M. S. The pricing of options and corporate liabilities. *Journal of Political Economy 81*, 3 (1973), 637–54.

[8] Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 31*, 3 (April 1986), 307–327.

[9] Bollerslev, T., Tauchen, G., and Zhou, H. Expected stock returns and variance risk premia. *Review of Financial Studies 22*, 11 (November 2009), 4463–4492.

[10] Bollerslev, T., and Todorov, V. Tails, fears, and risk premia. *The Journal of Finance 66*, 6 (2011), 2165–2211.

[11] Broadie, M., Johannes, M. S., and Chernov, M. Model Specification and Risk Premia: Evidence from Futures Options. *Journal of Finance 62*, 3 (2007), 1453–1490.

[12] Broadie, M., Johannes, M. S., and Chernov, M. Understanding Index Option Returns. *Review of Financial Studies 22*, 11 (2009), 4493–4529.

[13] Brownlees, C. T., Engle, R. F., and Kelly, B. T. A Practical Guide to Volatility Forecasting through Calm and Storm. *SSRN eLibrary* (2009).

[14] Brunnermeier, M. K., Gollier, C., and Parker, J. A. Optimal beliefs, asset prices, and the preference for skewed returns.

[15] Conrad, J., Dittmar, R., and Ghysels, E. Ex ante skewness and expected stock returns. *Journal of Finance* (forthcoming).

[16] Drechsler, I., and Yaron, A. What's vol got to do with it. *Review of Financial Studies 24*, 1 (2011), 1–45.

[17] Du, J., and Kapadia, N. The tail in the volatility index. Working paper, 2012.

[18] Engle, R. Risk and volatility: Econometric models and financial practice. *American Economic Review 94*, 3 (2004), 405–420.

[19] Engle, R. F. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica 50*, 4 (July 1982), 987–1007.

[20] Epstein, L. G., and Zin, S. E. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*, 4 (1989), pp. 937–969.

[21] Fama, E. F., and French, K. R. Multifactor explanations of asset pricing anomalies. *The Journal of Finance 51*, 1 (1996), 55–84.

[22] Farhi, E., Fraiberger, S. P., Gabaix, X., Ranciere, R., and Verdelhan, A. Crash risk in currency markets. NBER Working Papers 15062, National Bureau of Economic Research, Inc, June 2009.

[23] Gabaix, X. Power laws in economics and finance. *Annual Review of Economics 1*, 1 (2009), 255–294.

[24] Gabaix, X. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics 127*, 2 (2012), 645–700.

[25] Glosten, L. R., Jagannathan, R., and Runkle, D. E. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance 48*, 5 (December 1993), 1779–1801.

[26] Hill, B. M. A simple general approach to inference about the tail of a distribution. *The Annals of Statistics 3*, 5 (1975), pp. 1163–1174.

[27] Hodrick, R. Dividend yields and expected stock returns: alternative procedures for inference and measurement. *Review of Financial Studies 5*, 3 (1992), 357–386.

[28] Jurek, J., and Stafford, E. The cost of capital for alternative investments. *Harvard Business School Working Paper Series* (2012).

[29] Kelly, B. T. Tail Risk and Asset Prices. *Chicago Booth Working Paper* (2011).

[30] Kelly, B. T., Lustig, H. N., and Van Nieuwerburgh, S. Too-Systemic-To-Fail: What Option Markets Imply About Sector-Wide Government Guarantees. *Chicago Booth Working Paper* (2012).

[31] Nelson, D. B. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica 59*, 2 (1991), pp. 347–370.

[32] Newey, W. K., and West, K. D. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*, 3 (1987), pp. 703–708.

[33] OptionMetrics LLC. *Ivy DB File and Data Reference Manual*, September 2008.

[34] Pastor, L., and Stambaugh, R. F. Liquidity risk and expected stock returns. *Journal of Political Economy 111*, 3 (2003), pp. 642–685.

[35] Rietz, T. A. The equity risk premium: A solution. *Journal of Monetary Economics 22*, 1 (July 1988), 117–131.

[36] Santa-Clara, P., and Saretto, A. Option strategies: Good deals and margin calls. *Journal of Financial Markets 12*, 3 (2009), 391 – 417.

[37] Wachter, J. Can time-varying risk of rare disasters explain aggregate stock market volatility? Forthcoming, Journal of Finance, 2012.

[38] Wu, G. The determinants of asymmetric volatility. *Review of Financial Studies 14*, 3 (2001), 837–859.

[39] Xing, Y., Zhang, X., and Zhao, R. What does the individual option volatility smirk tell us about future equity returns? *Journal of Financial and Quantitative Analysis 45*, 03 (2010), 641–662.

A Appendix: Derivations

A.1 Useful Lemmas

Many derivations follow Farhi, et al. (2008), where I have adapted the math to fit the context of this paper. First, I will prove a very well-known Lemma which is the discrete time counterpart to Girsanov's Theorem:

Lemma 1. *Suppose (X, Y) are jointly Gaussian random variables under the physical probability measure $\mathbb{P}$. Define a new measure $\mathbb{Q}$ such that the Radon-Nikodym derivative is $d\mathbb{Q}/d\mathbb{P} = \exp(X - \mathbb{E}^{\mathbb{P}}[X] - var^{\mathbb{P}}(X)/2)$. Then under the new measure $\mathbb{Q}$, Y is Gaussian with the following distribution:*

$$Y \sim^{\mathbb{Q}} N(\mathbb{E}^{\mathbb{P}}[Y] + cov^{\mathbb{P}}(X, Y), var^{\mathbb{P}}(Y))$$

Proof. Calculate the moment generating function of Y under $\mathbb{Q}$ using the Radon-Nikodym derivative:

$$\begin{aligned}
\mathbb{E}^{\mathbb{Q}}[e^{kY}] &= \mathbb{E}^{\mathbb{P}}\left[\frac{d\mathbb{Q}}{d\mathbb{P}} e^{kY}\right] \\
&= \mathbb{E}^{\mathbb{P}}\left[e^{kY+X}\right] e^{-\mathbb{E}^{\mathbb{P}}[X] - var^{\mathbb{P}}(X)/2} \\
&= e^{k\,\mathbb{E}^{\mathbb{P}}[Y] + kcov^{\mathbb{P}}(Y,X) + k^2 var^{\mathbb{P}}(Y)/2}
\end{aligned}$$

where the last lines comes from the fact that $kY + X \sim N^{\mathbb{P}}(k\,\mathbb{E}^{\mathbb{P}}[Y] + \mathbb{E}^{\mathbb{P}}[X], 2kcov^{\mathbb{P}}(Y, X) + var^{\mathbb{P}}(X) + k^2 var^{\mathbb{P}}(Y))$. Thus, the moment generating function of Y under $\mathbb{Q}$ gives us the result. $\square$

Here is a useful Lemma that will be used to prove Theorem 1.

Lemma 2. *For $ln(A)$ Guassian distributed under $\mathbb{P}$:*

$$\begin{aligned}
E^{\mathbb{P}}[(C - A)^+] &= V^{call,BS}(C, \mathbb{E}^{\mathbb{P}}[A], var^{\mathbb{P}}(\ln(A))^{1/2}) \\
&= V^{put,BS}(\mathbb{E}^{\mathbb{P}}[A], C, var^{\mathbb{P}}(\ln(A))^{1/2})
\end{aligned}$$

where the convention is $V^{call,BS}(S, K, \sigma)$ is the Black-Scholes price of a call with spot rate S, strike K, interest rate 0 and horizon 1.

Proof. The Black-Scholes option pricing functions are defined as:

$$\begin{aligned}
V^{put,BS}(S, K, \sigma) &= \mathbb{E}^{\mathbb{Q}}[(K - Se^{\sigma u - \sigma^2/2})^+] \\
V^{call,BS}(S, K, \sigma) &= \mathbb{E}^{\mathbb{Q}}[(Se^{\sigma u - \sigma^2/2} - K)^+]
\end{aligned}$$

where $u \sim^{\mathbb{Q}} N(0, 1)$. The important thing to notice here is that $e^{\sigma u - \sigma^2/2}$ is log-normal with mean $-\sigma^2/2$ and variance σ^2 *under the measure* $\mathbb{Q}$, which is the same measure under which the expectation is taken over.

Under the assumption that $ln(A)$ is log-normal under measure $\mathbb{P}$, we can write $A = \mathbb{E}^{\mathbb{P}}[A]e^{X - var^{\mathbb{P}}(X)/2}$

where X is $\mathbb{P}$-Guassian with mean 0 and variance $var^{\mathbb{P}}(ln(A))$. Thus,

$$\begin{aligned}
E^{\mathbb{P}}[(C-A)^+] &= \mathbb{E}^{\mathbb{P}}[(C-\mathbb{E}^{\mathbb{P}}[A]e^{X-var^{\mathbb{P}}(X)/2})^+] \\
&= V^{put,BS}(\mathbb{E}^{\mathbb{P}}[A], C, var^{\mathbb{P}}(X)^{1/2}) \\
&= V^{put,BS}(\mathbb{E}^{\mathbb{P}}[A], C, var^{\mathbb{P}}(\ln(A))^{1/2})
\end{aligned}$$

which is trivial from the definition of the Black-Scholes put pricing formula. To get the result $E^{\mathbb{P}}[(C-A)^+] = V^{call,BS}(C, \mathbb{E}^{\mathbb{P}}[A], var^{\mathbb{P}}(\ln(A))^{1/2})$ we have to change measure using the Radon-Nikodym derivative $d\mathbb{Q}/d\mathbb{P} = \exp(X - var^{\mathbb{P}}(X)/2)$:

$$\begin{aligned}
E^{\mathbb{P}}[(C-A)^+] &= \mathbb{E}^{\mathbb{P}}[(C-\mathbb{E}^{\mathbb{P}}[A]e^{X-var^{\mathbb{P}}(X)/2})^+] \\
&= \mathbb{E}^{\mathbb{P}}[e^{X-var^{\mathbb{P}}(X)/2}(Ce^{Y}-\mathbb{E}^{\mathbb{P}}[A])^+] \\
&= \mathbb{E}^{\mathbb{P}}\left[\frac{d\mathbb{Q}}{d\mathbb{P}}(Ce^{Y}-\mathbb{E}^{\mathbb{P}}[A])^+\right] \\
&= \mathbb{E}^{\mathbb{Q}}[(Ce^{Y}-\mathbb{E}^{\mathbb{P}}[A])^+]
\end{aligned}$$

where $Y = -X + var^{\mathbb{P}}(X)/2$, which implies $\mathbb{E}^{\mathbb{P}}[Y] = var^{\mathbb{P}}(X)/2$ and $var^{\mathbb{P}}(Y) = var^{\mathbb{P}}(X)$. The challenge that remains is to figure out what the distribution of Y is under the measure $\mathbb{Q}$. Notice that $cov^{\mathbb{P}}(Y,X) = cov^{\mathbb{P}}(-X + var^{\mathbb{P}}(X)/2, X) = -var^{\mathbb{P}}(X)$. Additionally, (Y,X) is jointly normal under $\mathbb{P}$ by construction. Therefore, we can apply Lemma 1 to Y, X, and $d\mathbb{Q}/d\mathbb{P}$ and obtain $Y \sim^{\mathbb{Q}} N(-var^{\mathbb{P}}(X)/2, var^{\mathbb{P}}(X))$. In other words,

$$e^{Y} \sim^{\mathbb{Q}} e^{\sigma u - \sigma^2/2}$$

with $u \sim^{\mathbb{Q}} N(0,1)$ and $\sigma^2 = var^{\mathbb{P}}(X)$. Thus, applying the Black-Scholes call pricing formula (now under the measure $\mathbb{Q}$) we get:

$$\begin{aligned}
E^{\mathbb{P}}[(C-A)^+] &= \mathbb{E}^{\mathbb{Q}}[(Ce^{\sigma u-\sigma^2/2}-\mathbb{E}^{\mathbb{P}}[A])^+] \\
&= V^{call,BS}(C, \mathbb{E}^{\mathbb{P}}[A], var^{\mathbb{P}}(X)^{1/2}) \\
&= V^{call,BS}(C, \mathbb{E}^{\mathbb{P}}[A], var^{\mathbb{P}}(\ln(A))^{1/2})
\end{aligned}$$

□

A.2 Proof of Theorem 1

Now I will prove Theorem 1. Recall that Theorem 1 says that the measure of disaster risk (in the limit of small time) is $DR_{i,t}(M) = p_t\,\mathbb{E}_t[B_{t+1}^{-\gamma}(M - F_{t+1})]$.

Proof. Notice we can write the level of disaster risk as follows:

$$DR_{it}(M) = (1-p_t)DR_{it}^{ND}(M) + p_t DR_{it}^{D}(M) \tag{1}$$

where DR^D (resp DR^{ND} is the price of the disaster risk measure conditional on a disaster (resp.

no disaster) happening next period. Let us focus on the value of DR_{it} conditional on no disaster:

$$
\begin{aligned}
DR_{it}^{ND}(M) &= e^{-\delta}\left(\mathbb{E}_t^{ND}[(M - P_{t+1}/P_t)^+] - Me^{-\mu_{it}}\,\mathbb{E}_t^{ND}[(P_{t+1}/P_t - M^{-1}e^{2\mu_{it}})^+\right) \\
&= e^{-\delta}\left(V^{call,BS}(M, e^{\mu_{it}}, \sigma) - Me^{-\mu_{it}}V^{call,BS}(e^{\mu_{it}}, M^{-1}e^{2\mu_{it}}, \sigma)\right) \\
&= e^{-\delta}\left(Me^{-\mu_{it}}V^{call,BS}(e^{\mu_{it}}, M^{-1}e^{2\mu_{it}}, \sigma) - Me^{-\mu_{it}}V^{call,BS}(e^{\mu_{it}}, M^{-1}e^{2\mu_{it}}, \sigma)\right) \\
&= 0 \qquad (2)
\end{aligned}
$$

where the second line is by Lemma 2 and the definition of the Black-Scholes pricing formula used in Lemma 2. The third line is using the homogeneity of degree 1 (in both the first and second arguments) of the Black-Scholes formula.

Therefore, the measure of disaster risk comes entirely from the disaster component of the put and call options:

$$
\begin{aligned}
DR_{it}(M) &= p_t DR_{it}^{D}(M) \\
&= e^{-\delta}p_t\left(\mathbb{E}_t^{D}[B_{t+1}^{-\gamma}(M - P_{t+1}/P_t)^+] - Me^{-\mu_{it}}\,\mathbb{E}_t^{D}[B_{t+1}^{-\gamma}(P_{t+1}/P_t - M^{-1}e^{2\mu_{it}})^+\right) \\
&= e^{-\delta}p_t\left(\mathbb{E}_t^{D}[B_{t+1}^{-\gamma}(M - e^{\mu_{it}}F_{i,t+1})^+] - Me^{-\mu_{it}}\,\mathbb{E}_t^{D}[B_{t+1}^{-\gamma}(e^{\mu_{it}}F_{i,t+1} - M^{-1}e^{2\mu_{it}})^+\right) \\
&= e^{-\delta}p_t\left(\mathbb{E}_t^{D}[B_{t+1}^{-\gamma}(M - e^{\mu_{it}}F_{i,t+1})^+] - \mathbb{E}_t^{D}[B_{t+1}^{-\gamma}(MF_{i,t+1} - e^{\mu_{it}})^+\right) \qquad (3)
\end{aligned}
$$

Now, if we take a small-time limit the non-stochastic portion of the stochastic discount factor growth rate and the non-stochastic portion of the stock growth rate go to 1, which reduces Equation 3 to:

$$
\begin{aligned}
DR_{it}(M) &= p_t\left(\mathbb{E}_t^{D}[B_{t+1}^{-\gamma}(M - F_{i,t+1})^+] - \mathbb{E}_t^{D}[B_{t+1}^{-\gamma}(MF_{i,t+1} - 1)^+\right) \\
&= p_t\,\mathbb{E}_t^{D}[B_{t+1}^{-\gamma}(M - F_{i,t+1})]
\end{aligned}
$$

where the last line comes from the assumptions of $M > F_{i,t+1}$ *a.s* and $MF_{i,t+1} < 1$ *a.s.*. It is important to note that, contrary to many studies which study jumps in the return process, taking a small time limit does not mean the value of the option comes entirely from the "jump" component. This is intuitively because I am dealing with *consumption* disasters (or jumps) and thus the probability of a disaster over a small time interval does not go to 1. Put differently, I am assuming that over a small time interval (here a month) there is a small probability of a jump and a large probability of log-normality.

□

A.3 Proof of Theorem 2

The proof follows naturally from the proof of Theorem 1, with the only difference being we now use an sdf in an Epstein-Zin world. To start, I will show how equation (3) still holds in an Epstein-Zin world. The price process in an Epstein-Zin world is given by (13), and can be compactly expressed as:

$$
\begin{aligned}
\left(\frac{P_{i,t+1}}{P_{it}}\right) &= \left(\frac{D_{i,t+1}}{D_{it}}\right) \times \left(\frac{\tilde{a} + \tilde{b}\hat{H}_{i,t+1}^{EZ}}{\tilde{a} + \tilde{b}\hat{H}_{i,t}^{EZ}}\right) \\
&= \left(\frac{D_{i,t+1}}{D_{it}}\right) \times \left(\frac{\tilde{a} + \tilde{b}\left[\frac{1+H_{i*}^{EZ}}{1+H_{it}^{EZ}} e^{-\phi_H}\hat{H}_{it}^{EZ} + \epsilon_{i,t+1}^{EZ,H}\right]}{\tilde{a} + \tilde{b}\hat{H}_{i,t}^{EZ}}\right)
\end{aligned}
$$

for some constants $\tilde{a}, \tilde{b}$. Thus, the sources of uncertainty in log-returns come from the dividend growth process, $\epsilon_{i,t+1}^{D}, F_{i,t+1}$ and the innovation to EZ-enriched reslience, $\epsilon_{i,t+1}^{EZ,H}$. Analogously, it is innovations to the dividend process and the power-utility reslience that drive the return process innovations in the power-utility version of the model. Hence, to generate a return process that is consistent with (13) and (3) in the EZ world, I will rely on the same argument that was used in the power-utility case. That is, define:

$$
\begin{aligned}
Y &= \left(\frac{\tilde{a} + \tilde{b}\left[\frac{1+H_{i*}^{EZ}}{1+H_{it}^{EZ}} e^{-\phi_H}\hat{H}_{it}^{EZ} + \epsilon_{i,t+1}^{EZ,H}\right]}{\tilde{a} + \tilde{b}\hat{H}_{i,t}^{EZ}}\right) e^{g_{iD}} \\
Z &= e^{\mu + \sigma u_{t+1} - \sigma^2/2} \\
X &= 1 + \epsilon_{i,t+1}^{D}
\end{aligned}
$$

Applying Lemma 3 from the Online Appendix to Gabaix (2012) and defining $(1 + \epsilon_{i,t+1}^{D})$ as in the power utility case then delivers a price process that is still consistent with equation (3).

All I have done thus far is to square the specification that returns are log-normal in times of no disasters and fall by $F_{i,t+1}$ in times of disasters with the Epstein-Zin extension of the model. Now I will re-derive the measure of disaster risk in this world. Theorem 3 of Gabaix (2012) states thats that, to a leading order, the stochastic discount factor in an Epstein-Zin world is:

$$
\frac{M_{t+1}}{M_t} = e^{-\delta}(1 + (\chi - 1)H_{Ct} + \epsilon_{t+1}^{M}) \times \begin{cases} 1 & \text{if there is no disaster at } t+1 \\ B_{t+1}^{-\gamma} & \text{if there is a disaster at } t+1 \end{cases}
$$

with $\delta = \rho + g_c/\chi$ and ϵ_{t+1}^{M} a linear function of the shock to consumption resilience (so it is independent of other shocks in the economy as well and still has a mean of zero). Define $\eta_t = 1 + (\chi - 1)H_{Ct}$. To start, let's focus on the value of a disaster risk given there is no disaster:

$$
\begin{aligned}
DR_{it}^{ND}(K) &= e^{-\delta}\big(\mathbb{E}_t^{ND}[(\eta_t + \epsilon_{t+1}^{M})(K - P_{t+1}/P_t)^{+}] \\
&\quad - Ke^{-\mu_{it}}\,\mathbb{E}_t^{ND}[(\eta_t + \epsilon_{t+1}^{M})(P_{t+1}/P_t - K^{-1}e^{2\mu_{it}})^{+}\big) \\
&= e^{-\delta}\eta_t\left(V^{call,BS}(K, e^{\mu_{it}}, \sigma) - Ke^{-\mu_{it}}V^{call,BS}(e^{\mu_{it}}, K^{-1}e^{2\mu_{it}}, \sigma)\right) \\
&= e^{-\delta}\eta_t\left(Ke^{-\mu_{it}}V^{call,BS}(e^{\mu_{it}}, K^{-1}e^{2\mu_{it}}, \sigma) - Ke^{-\mu_{it}}V^{call,BS}(e^{\mu_{it}}, K^{-1}e^{2\mu_{it}}, \sigma)\right) \\
&= 0
\end{aligned}
$$

where the second line is by Lemma 2 and the definition of the Black-Scholes pricing formula used in Lemma 2. The third line is using the homogeneity of degree 1 (in both the first and second arguments) of the Black-Scholes formula.

As in the power utility case, the measure of disaster risk comes entirely from the disaster component of the put and the call:

$$
\begin{aligned}
DR_{it}(K) &= p_t DR_{it}^D(K) \\
&= e^{-\delta} p_t \big(\mathbb{E}_t^D[(\eta_t + \epsilon_{t+1}^M) B_{t+1}^{-\gamma}(K - P_{t+1}/P_t)^+] \\
&\quad - K e^{-\mu_{it}} \,\mathbb{E}_t^D[(\eta_t + \epsilon_{t+1}^M) B_{t+1}^{-\gamma}(P_{t+1}/P_t - K^{-1}e^{2\mu_{it}})^+\big) \\
&= e^{-\delta} p_t \eta_t \Big(\mathbb{E}_t^D[B_{t+1}^{-\gamma}(K - e^{\mu_{it}} F_{i,t+1})^+] - K e^{-\mu_{it}} \,\mathbb{E}_t^D[B_{t+1}^{-\gamma}(e^{\mu_{it}} F_{i,t+1} - K^{-1} e^{2\mu_{it}})^+\Big) \\
&= e^{-\delta} p_t \eta_t \Big(\mathbb{E}_t^D[B_{t+1}^{-\gamma}(K - e^{\mu_{it}} F_{i,t+1})^+] - \mathbb{E}_t^D[B_{t+1}^{-\gamma}(K F_{i,t+1} - e^{\mu_{it}})^+\Big) \qquad (4)
\end{aligned}
$$

Now, if we take a small-time limit the non-stochastic portion of the stochastic discount factor growth rate and the non-stochastic portion of the stock growth rate go to 1, which reduces Equation 4 to:

$$
\begin{aligned}
DR_{it}(K) &= p_t \eta_t \Big(\mathbb{E}_t^D[B_{t+1}^{-\gamma}(K - F_{i,t+1})^+] - \mathbb{E}_t^D[B_{t+1}^{-\gamma}(K F_{i,t+1} - 1)^+\Big) \\
&= p_t \eta_t \,\mathbb{E}_t^D[B_{t+1}^{-\gamma}(K - F_{i,t+1})]
\end{aligned}
$$

where the last line comes from the assumptions of $K > F_{i,t+1}$ $a.s$ and $KF_{i,t+1} < 1$ $a.s.$. Thus, the DR_{it}l in the Epstein-Zin is the power-utility version multiplied by the factor η_t. Recall $\eta_t = 1 + (\chi - 1)H_{Ct} \approx 1 + (\chi - 1)H_{C*}$, where $H_{C*} = p^*[B^{1-\gamma} - 1]$. Using the calibration values in Table 6, η_t will behave like an AR(1) process that hovers around its long-run mean of approximately 0.92. Moreover, the mean-reversion and volatility of this process will mirror the resilience of aggregate consumption, and thus will have very little volatility and will be a very slow moving process. Therefore, η_t will behave much like a constant compared to the other dynamic variables that comprise my disaster risk measure and I conclude it is reasonable to assume $\eta_t \approx 1$.

B Appendix: Empirical Methodology

B.1 General Data and Construction Details

The generic dataset that I use for this paper is the OptionsMetrics Volatility Surface from WRDS. Specifically, I examine 30-day options on S&P 500 firms. To be specific, data for a given month m is obtained as follows:

1. Using Compustat, I obtain the members of the S&P 500 Composite Index as of the end of month m.[29] This is critical for avoiding a lookahead bias when forming the disaster risk equity portfolios.

2. When estimating the probability of a disaster, this is not a necessary consideration and hence I use all firms in the master data set (across all months) for this task only.

3. Once I obtain the list of firms that were in the index at month m, I download their OMVS history for month m. Often options are not available for all tickers so I obtain the maximum amount possible

4. The OMVS history for month m forms the base dataset for the rest of the analysis, and I repeat the procedure each month

The filtering process thereafter is simple: within a month I exclude a firm if it has more than 10 observations missing. Finally, there are some cases in the WRDS database where a firm (which is uniquely identified by "secid") has large gaps in its data series. These gaps correspond to months at a time in some cases. To avoid any issues with these gaps, I exclude firms that have more than a month missing in the volatility surface or price series. In the end, this leaves me with, on average, roughly 350 firms per month. The full list of firms considered in each portfolio on a monthly basis is available upon request. The data used is downloaded from WRDS and read into an SQL server I maintain for personal use. The code used to construct all series in the paper is also available upon request, and operates on this SQL server.

B.2 Disaster Risk from the OptionsMetrics Volatility Surface

B.2.1 Methodology for DR_{it}

In theory, the construction of a disaster risk is straightforward and involves the following steps:

1. Pick a desired moneyness level M

2. Calculate the value of a portfolio that is long 1 unit of a put at moneyness M, and short M-units of a call at moneyness M^{-1}. The two restrictions are that $MF_{i,t+1} < 1$ and $M < F_{i,t+1}$. In practice, these two restrictions will be satisfied since I am not going to purchase very deep OTM put options.

It is worth noting again that since the OMVS implied volatility surface is interpolated each day, there is an interpolated set of 30-day options each day of the sample so term structure issues are non-existent in this context.[30] The two main issues one must consider are 1) which moneyness M to choose, and 2) a call of moneyness M^{-1} is not typically available on this interpolated surface.

[29] Index ticker "i0003"

[30] Though it would be interesting to consider a term structure of tail risk.

Moneyness versus Delta

To address the first issue, I instead buy the put option which has a Black-Scholes delta of -25. This corresponds to an OTM put, whose moneyness varies through time with the implied volatility of the option. In fact, OptionMetrics constructs its volatility surface by interpolating over Black-Scholes delta (as opposed to strike). Thus, keeping the delta of DR_{it} constant through time takes into account changes in volatility in a natural way since the delta adjusts for this. I want to emphasize that this is does not bind me to a Black-Scholes world for asset prices. The Black-Scholes delta is simply a function that links moneyness to implied volatility. I could very well have taken, for example, the moneyness with a strike that is one-standard deviation away from the current stock price. The reason for using the Black-Scholes delta is simply because that is how the data is packaged.

OptionsMetrics also reports the implied strike for each of its interpolated options, which is calculated by inverting the standard Black-Scholes delta formula. Thus, moneyness in this context can be thought of as a function of the reported Black-Scholes delta (recall the surface is interpolated over delta):

$$\begin{aligned} M_t^{implied} &= \frac{K_t^{implied}}{S_t} \\ &= \exp\left(\Phi^{-1}(-\Delta_{BS}^P)\sigma_t^{P,IV}\sqrt{\tau} + (r_t - q_t + (\sigma_t^{P,IV})^2/2)\tau\right) \\ &= f(\Delta_{BS}^P; \sigma_t^{P,IV}, \tau, r_t, q_t, S_t) \end{aligned}$$

where $\Phi(\cdot)$ is the Normal cdf, Δ_{BS}^P is the Black-Scholes delta of the put, τ is the time to maturity, r_t is the risk-free rate on date t, q_t is the dividend yield on date t, and $\sigma_t^{P,IV}$ is the implied volatility of the associated put option on date t. Since $\Delta_{BS}^P = -25$ and $\tau = 1/12$ are held constant over time, the time variation in the moneyness of the put comes primarily from time variation in implied volatility.

Once I have the implied moneyness of the -25 delta put, the next task is to calculate the price of the $(M_t^{implied})^{-1}$ call option. There are two alternative ways to do so:

1. Using the implied strikes and implied volatilities from OMVS, interpolate over this surface to find the implied volatility of the desired call option. Then use the Black-Scholes formula with associated known inputs to calculate the price of the call option needed for the disaster risk measure.

2. Using the call option with $\Delta_{BS}^C = -\Delta_{BS}^P$, which is already reported on the OMVS.

The drawback of the first approach is a lot of estimation error, since I am interpolating over an already interpolated surface. The drawback of the second approach is that the call with $\Delta_{BS}^C = -\Delta_{BS}^P$ is not necessarily the call option with moneyness $(M_t^{implied})^{-1}$. However, in Appendix B.2.4 I show that in small time intervals an option with $\Delta_{BS}^C = -\Delta_{BS}^P$ is in fact the option with moneyness $(M_t^{implied})^{-1}$. Moreover, I derive the approximation error as a function of the slope of the implied volatility surface.[31] Both approaches – interpolating over the OMVS surface vs. using the call with $\Delta_{BS}^C = -\Delta_{BS}^P$ – have their own problems, but I choose to match the deltas of the put and call since these options are already reported by OMVS.[32]

[31]Intuitively, the approximation contains the most error when the implied volatility is the most steep, which is likely in a crisis. Still, the OMVS surface itself is quite noisy (as measured by their "dispersion" variable) in these periods so it is unclear whether interpolating over the OMVS presents a better alternative.

[32]I tried both approaches and the "double interpolation" method is indeed much less reliable.

In light of the previous discussion, I will quote my disaster risk measure in terms of delta, with the implicit understanding that a "25-Δ" DR_{it} implies buying the put with $\Delta_{BS}^{P} = -25$ and selling the call with $\Delta_{BS}^{C} = 25$. The number of units of the call sold is calculated using the reported implied strike for the -25-delta put. To summarize, my empirical approach for constructing my disaster risk measure at a given Δ is:

1. On each day, choose the put option with $\Delta_{BS}^{P} = -\Delta$ and find the associated price
2. Using the implied strike and the spot price, calculate the moneyess of the option
3. Choose the call option with $\Delta_{BS}^{C} = \Delta$ to complete the measure of disaster risk

My disaster risk measure quoted in terms of delta is thus,

$$\begin{aligned} DR_{i,t}(\Delta) &= Put(\Delta_{BS}^{P} = -\Delta) - M_t^{implied,put} \cdot Call(\Delta_{BS}^{C} = \Delta) \\ M_t^{implied,put} &= \exp\left(\Phi^{-1}(-\Delta)\sigma_t^{P,IV}\sqrt{\tau} + (r_t - q_t + (\sigma_t^{P,IV})^2/2)\tau\right) \end{aligned} \tag{1}$$

B.2.2 Methodology for DR_{it} Implied Equity Premium and Disaster Probability

Corollary 1 shows how the DR_{it} implied equity premium can be derived using the average of two measures whose moneyness adds up to 2. In terms of delta, this approximately corresponds to two $DR_{it}(M)$ measures whose deltas add up to 1.[33] For example, one could back out the implied equity premium using the average of my measure with $\Delta = 40$ and $\Delta = 60$. The OptionMetrics surface is interpolated over deltas (in absolute value) of $|\Delta| = \{20, 25, 30, ..., 70, 75, 80\}$. Thus, I come up with an estimate of the implied equity premium using each of the $\{45, 55\}, \{40, 60\}, \{35, 65\}, \{30, 70\}, \{25, 75\}, \{20, 80\}$ pairs DR_{it}. Finally, I take an average of these individual estimates to deliver the implied equity premium on a given day.

From Corollary 2 we can back out the risk-neutral probability of disaster using the difference between two disaster risk measures. I chose to use the disaster risk measure with $\Delta = 25$ and also the DR_{it} with $\Delta = 20$. Naturally, the DR_{it} with $\Delta = 25$ is closer to the money and thus will have a higher strike for the put portion of the measure.

B.2.3 Matching Deltas versus Symmetrical Calls/Puts

In this section, I will argue that matching the delta of a put and a call to construct my DR measure is a valid approximation to matching a symmetrical put and call.[34] Furthermore, I will derive the approximation error of matching deltas, since the theory suggests this measure should be constructed using symmetrical puts and calls. I will abuse notation by dropping time subscripts, with the implicit understanding that my DR measure is constructed each day.

Consider a put option with a strike of $K^{put} = Se^{\theta_p\sqrt{\tau}} \Rightarrow m_{put} \equiv K^{put}/S = e^{\theta_p\sqrt{\tau}}$. Futhermore, denote $\sigma_{IV}^{\Delta_{BS}^{P}}$ as implied volatility of an option with a delta equal to Δ_{BS}^{P}. The reason for this rather

[33] This can be easily seen in the data.
[34] By symmetrical I mean a put and a call whose moneyness are inverses.

cumbersome notation will become clear. Now consider the Black-Scholes delta of this option:

$$\Delta_{BS}^{P} = -e^{-q\tau}\Phi\left(-\frac{\ln(S/K) + (r - q + (\sigma_{IV}^{\Delta_{BS}^{P}})^2/2)\tau}{\sigma_{IV}^{\Delta_{BS}^{P}}\sqrt{\tau}}\right)$$
$$= -e^{-q\tau}\Phi\left(-\frac{-\theta_p\sqrt{\tau} + (r - q + (\sigma_{IV}^{\Delta_{BS}^{P}})^2/2)\tau}{\sigma_{IV}^{\Delta_{BS}^{P}}\sqrt{\tau}}\right)$$

Taking the limit as $\tau \searrow 0$:

$$\lim_{\tau\searrow 0} \Delta_{BS}^{P} = -\Phi\left(\frac{\theta_p}{\sigma_{IV}^{\Delta_{BS}^{P}}}\right)$$

Thus, for small-time limits, we can express θ_p in terms of Black-Scholes delta as:

$$\theta_p = \Phi^{-1}\left(-\Delta_{BS}^{P}\right)\cdot\sigma_{IV}^{\Delta_{BS}^{P}} \tag{2}$$

Now consider the same exercise for a call option. That is, let $K^{call} = Se^{\theta_c\sqrt{\tau}} \Rightarrow m_{call} \equiv K^{call}/S = e^{\theta_c\sqrt{\tau}}$. Notice that $m_{call} * m_{put} = 1$ implies $\theta_p = -\theta_c$.[35] Thus, in order to show that matching deltas is indeed the same as matching symmetrical put and calls, I simply need to show that the matching deltas implies $\theta_p = -\theta_c$. The expression for θ_c as a function of Black–Scholes delta is derived in a similar fashion:

$$\lim_{\tau\searrow 0} \Delta_{BS}^{C} = \Phi\left(\frac{-\theta_c}{\sigma_{IV}^{\Delta_{BS}^{C}}}\right)$$

Inverting this expression gives us:

$$\theta_c = -\Phi^{-1}\left(-\Delta_{BS}^{C}\right)\cdot\sigma_{IV}^{\Delta_{BS}^{C}} \tag{3}$$

Thus, if $\sigma_{IV}^{\Delta_{BS}^{C}} = \sigma_{IV}^{\Delta_{BS}^{P}}$ then setting $\Delta_{BS}^{C} = -\Delta_{BS}^{P}$ indeed implies $\theta_p = -\theta_c$, which gives the desired result.

B.2.4 Error in Matching Deltas versus Interpolating

Clearly, the derivation in the previous subsection relies on a flat implied volatility curve and also that even in a small time limit the strike can be expressed in the form $K = Se^{\theta\sqrt{\tau}}$. The former assumption is obviously false, since it is well known that the implied volatility curve is in fact not flat. The latter assumption seems more reasonable since I am dealing with options with $\tau = 1/12$ and τ goes to zero faster than $\sqrt{\tau}$ so the limiting arguments in that dimension likely still hold.

I will now derive the approximation error from matching delta for DR_{it}. I will do this in terms of how far apart the implied moneyness of delta matched calls and puts are different from being

[35] A subtle point is that τ must be near zero but not zero exactly.

reciprocals. Repeating the formula for Black-Scholes delta in terms of $m_{put} = K_{put}/S$:

$$\Delta_{BS}^{P} = -e^{-q\tau} N\left(-\frac{-\ln(m_{put}) + (r - q + (\sigma_{IV}^{\Delta_{BS}^{P}})^2/2)\tau}{\sigma_{IV}^{\Delta_{BS}^{P}}\sqrt{\tau}}\right)$$

Inverting this gives us moneyness in terms of delta, which is how the OptionMetrics data comes packaged:

$$m_{put} = \exp\left(\Phi^{-1}\left(-e^{q\tau}\Delta_{BS}^{P}\right)\sigma_{IV}^{\Delta_{BS}^{P}}\sqrt{\tau} + (r - q + (\sigma_{IV}^{\Delta_{BS}^{P}})^2/2)\tau\right) \tag{4}$$

Recall that DR_{it}l consists of buying a call whose moneyness is given by $m_{call}^{target} \equiv m_{put}^{-1}$:

$$\begin{aligned} m_{call}^{target} &\equiv m_{put}^{-1} \\ &= \exp\left(-\Phi^{-1}\left(-e^{q\tau}\Delta_{BS}^{P}\right)\sigma_{IV}^{\Delta_{BS}^{P}}\sqrt{\tau} - (r - q + (\sigma_{IV}^{\Delta_{BS}^{P}})^2/2)\tau\right) \end{aligned} \tag{5}$$

Using the same logic, we can write the moneyness of a call in terms of its Black–Scholes delta:

$$\begin{aligned} m_{call} &= \exp\left(-\Phi^{-1}\left(e^{q\tau}\Delta_{BS}^{C}\right)\sigma_{IV}^{\Delta_{BS}^{C}}\sqrt{\tau} + (r - q + (\sigma_{IV}^{\Delta_{BS}^{C}})^2/2)\tau\right) \\ &= \exp\left(-\Phi^{-1}\left(e^{q\tau}\Delta_{BS}^{C}\right)\sigma_{IV}^{\Delta_{BS}^{C}}\sqrt{\tau} + (r - q + (\sigma_{IV}^{\Delta_{BS}^{C}})^2/2)\tau\right) \end{aligned}$$

When I match deltas I set $\Delta_{BS}^{C} = -\Delta_{BS}^{P}$. Thus, the effect of delta matching means the moneyness of the purchased call will be:

$$m_{call} = \exp\left(-\Phi^{-1}\left(-e^{q\tau}\Delta_{BS}^{P}\right)\sigma_{IV}^{\Delta_{BS}^{C}}\sqrt{\tau} + (r - q + (\sigma_{IV}^{\Delta_{BS}^{C}})^2/2)\tau\right) \tag{6}$$

Thus, the approximation error from delta matching will precisely be the difference between selling a call at m_{call}^{target} versus selling a call at m_{call}, or rather the difference between expressions (5) and (6). To quantify this, we can compare the log-moneyness of the target call and the log-moneyness of the delta-matched call. For further simplification, let $r = q = 0$ so the only differences arise from the implied volatility surface. Finally let $\eta \equiv -\Phi^{-1}\left(\Delta_{BS}^{P}\right) > 0$ for OTM puts. Then Equations (5) and (6) reduce to:

$$\begin{aligned} \log(m_{call}^{target}) &= \eta\sigma_{IV}^{\Delta_{BS}^{P}}\sqrt{\tau} - (\sigma_{IV}^{\Delta_{BS}^{P}})^2 \cdot \tau/2 \\ \log(m_{call}) &= \eta\sigma_{IV}^{\Delta_{BS}^{C}}\sqrt{\tau} + (\sigma_{IV}^{\Delta_{BS}^{C}})^2 \cdot \tau/2 \end{aligned} \tag{7}$$

In practice, the slope of the implied volatility surface is typically negative. That is, OTM puts have higher implied volatilities than OTM call options. Since I am concerned with OTM options (or equivalently, options whose delta is less than 0.5 in absolute value), a reasonable assumption to make is:[36]

$$\kappa \equiv \sigma_{IV}^{\Delta_{BS}^{P}} - \sigma_{IV}^{\Delta_{BS}^{C}} > 0$$

[36] Note that this will be the difference in implied volatilities for a put option and a call option whose Black-Scholes deltas are the same in absolute value. For example, for the baseline DR measure I use deltas of 25, so κ is the difference between the implied volatility of a put with delta of -25 and a call of delta with 25.

Using the identity $\sigma_{IV}^{\Delta_{BS}^P} = \sigma_{IV}^{\Delta_{BS}^C} + \kappa$, some algebra shows:

$$\log(m_{call}^{target}) = \log(m_{call}) + \left(\eta\kappa\sqrt{\tau} - [(\sigma_{IV}^{\Delta_{BS}^C})^2 + \kappa\sigma_{IV}^{\Delta_{BS}^C} + \kappa^2/2]\tau\right) \tag{8}$$

So the second term on the RHS is the error in moneyness from matching deltas instead of matching moneyness directly. Clearly, as τ gets small this approximation error will also disappear, but of course this will depend on the slope of the volatility surface (roughly κ) and the implied volatility of the delta-matched call. Note further that this is an error in moneyness, but when we apply the Black-Scholes function which is smooth and nearly flat for very OTM call options, the error in moneyness will be further diluted.

Notice the entire error derivation relies on being able to calculate the option price of a call with m_{call}^{target}. However, in practice this option is not available on the OMVS surface, so calculating its associated implied volatility (and hence call price) requires further interpolation over an already interpolated surface. I take the stance that the error in this process is less than the error in Equation (8); hence, I argue that matching the deltas is, in a sense, the lesser of two evils.

C Appendix: Primer on Linearity-Generating Processes

In order to keep the paper (relatively) self-contained, I will present a short primer on the linearity-generating processes used to derive closed form solutions for the price-dividend ratios in this paper. The seminal paper on LG-processes is Gabaix (2009) and is highly recommended for a thorough understanding of LG-processes. An LG process is simply a stochastic process with the very special property that it yields linear expressions for stocks and bonds, with an arbitrary number of factors. As all of asset pricing relies on stochastic discount factors, the LG process is defined through two main moment conditions, which are easily interpreted in a stochastic discount factor framework. I start by repeating Definition 2 of Gabaix (2009),

Definition 2. *The stochastic process $M_t D_t (1, X_t')_{t=0,1,\ldots}$, with $M_t D_t \in \mathbb{R}^+$ and $X_t \in \mathbb{R}^n$ is a linearity-generating process if it is in L^1 and there are constants $\alpha \in \mathbb{R}, \gamma, \delta \in \mathbb{R}^n, \Gamma \in \mathbb{R}^{n\times n}$, such that the following moment conditions hold at all $t \in \mathbb{N}$:*

$$\mathbb{E}_t \left[\frac{M_{t+1} D_{t+1}}{M_t D_t} \right] = \alpha + \delta' X_t, \tag{9}$$

$$\mathbb{E}_t \left[\frac{M_{t+1} D_{t+1}}{M_t D_t} X_{t+1} \right] = \gamma + \Gamma X_t \tag{10}$$

Here we can think of X_t as a set of factors that determine dividend growth and/or the stochastic discount factor. M_t can be thought of as a candidate SDF, and D_{t+1} is a dividend process. For ease of exposition, I will focus on the case where $n = 1$. Equation (9) simply states that the risk-adjusted expected dividend growth is linear in the factor X_t. Equation (10) is a statement about the process X_{t+1}. Consider the case where $M_t = 1, D_t = 1, \forall t$: then Equation (10) says that X_t is an AR(1) process. Hence, Gabaix (2009) describes X_t as following a "twisted" AR(1). It turns out that the "twist" here allows the future discounted stream of cash-flows to be a linear function of X_t as opposed to, for example, an infinite sum as in the affine class of models from Duffie and Kan (1996). Theorem 2 in Gabaix (2009) shows that if $M_t D_t (1, X_t')$ satisfies Definition 2, then:

Theorem (Theorem 2 from Gabaix (2009))**.** *Define:*

$$\Omega = \begin{pmatrix} \alpha & \delta' \\ \gamma & \Gamma \end{pmatrix}$$

If P_t is the value of a stock that pays dividend D_t at time t, M_t is the stochastic discount factor, and the $|\Omega| < 1$, then the price-dividend ratio of the stock is:

$$\begin{aligned} \frac{P_t}{D_t} &= \mathbb{E}_t \left[\sum_{s=t+1}^{\infty} \frac{M_s D_s}{M_t D_t} \right] \\ &= (1 \quad 0_n) \Omega (I_{n+1} - \Omega)^{-1} (1 \quad X_t)' \end{aligned} \tag{11}$$

This theorem demonstrates the usefulness of the LG class, namely that it delivers price-dividend ratios that are linear in arbitrary state-variables. In practice, we can generically satisfy (9), and then we reverse engineer a process for the state-variable that will allow (10) to hold. For example, consider the power-utility version of the model in this paper. Here,

$$\frac{M_{t+1}}{M_t} = e^{-\delta} \times \begin{cases} 1 & \text{if there is no disaster at } t+1 \\ B_{t+1}^{-\gamma} & \text{if there is a disaster at } t+1 \end{cases}$$

Let's look at condition (9):

$$
\begin{aligned}
\frac{M_{t+1}D_{t+1}}{M_tD_t} &= e^{-\delta+g_{iD}}(1+\epsilon^D_{t+1}) \times \begin{cases} 1 & \text{if there is no disaster at } t+1 \\ B^{-\gamma}_{t+1}F_{i,t+1} & \text{if there is a disaster at } t+1 \end{cases} \\
\mathbb{E}_t\left[\frac{M_{t+1}D_{t+1}}{M_tD_t}\right] &= e^{-\delta+g_{iD}}\left\{(1-p_t)\cdot 1 + p_t\,\mathbb{E}_t[B^{-\gamma}_{t+1}F_{i,t+1}]\right\} \\
&= e^{\delta+g_{iD}}\left\{1+\underbrace{p_t\,\mathbb{E}_t[B^{-\gamma}_{t+1}F_{i,t+1}-1]}_{\equiv H_{it}}\right\} \\
&= \alpha + \delta H_{it}
\end{aligned}
$$

Now it is easy to see where the definition of resilience came from: $H_{it} = p_t\,\mathbb{E}_t[B^{-\gamma}_{t+1}F_{i,t+1}-1]$. We can think of this as our "state" variable. Following with the definition of LG-processes, we will need to design this as a "twisted" AR(1) process in order to apply Theorem 2 of Gabaix (2009). Let's naturally define the level to which it mean reverts (it's typical value) as H_{i*}. It is then easier to decompose H_{it} into a constant part and a variable part that hovers around zero (thereby making the entire process hover around H_{i*}). That is, let $H_{it} = H_{i*} + \hat{H}_{it}$.[37] Our goal is to have the price-dividend ratio be a linear function of $\hat{H}_{it}$, so we need to reverse engineer a process for $\hat{H}_{it}$ that will satisfy condition (10). $\hat{H}_{it}$ will behave like an AR(1) with an unconditional mean of 0, so define its mean-zero shocks as ϵ^H_{t+1} and assume they are independent of any other variable. We can generically write:

$$
\hat{H}_{i,t+1} = f(\hat{H}_{it}) + \epsilon^H_{t+1}
$$

$$
\begin{aligned}
\mathbb{E}_t\left[\frac{M_{t+1}D_{t+1}}{M_tD_t}\hat{H}_{it+1}\right] &= \mathbb{E}_t\left[\frac{M_{t+1}D_{t+1}}{M_tD_t}\right]f(\hat{H}_{it}) \\
&= e^{-\delta+g_{iD}}[1+H_{it}]f(\hat{H}_{it})
\end{aligned}
$$

The RHS must be linear in $\hat{H}_{it}$ in order for us to apply Theorem 2 of Gabaix (2009). We can then deduce that $f(\cdot)$ must have a term $(1+H_{it})$ in the denominator. By defining:

$$
f(\hat{H}_{it}) = \frac{1+H_{i*}}{1+H_{it}}e^{-\phi_H}\hat{H}_{it}
$$

we can satisfy the condition (10). Thus, Theorem 2 of Gabaix (2009) delivers us price-dividend ratios that are linear in $\hat{H}_{it}$ and the constants are derived from the moment conditions (9) and (10).

Some notes are in order concerning the functional form $f(\cdot)$. First, the exponent $e^{-\phi_H}$ simply lets the modeler choose the speed of mean-reversion (to zero) for the process $\hat{H}_{it}$. To this end, the constant $1+H_{i*}$ in the denominator allows ϕ_H alone to control the speed of mean-reversion since to a leading order then $\hat{H}_{i,t+1} \approx e^{-\phi_H}\hat{H}_{it} + \epsilon^H_{t+1}$. To summarize:

1. Once the primitives of the model deliver M_{t+1}, we can solve for (9).
2. The definition of resilience follows naturally from the expression in (9).
3. Design a process for resilience such that (10) holds.

[37] In reality, then $\hat{H}_{it}$ is the state variable instead of H_{it}.

4. Use Theorem 2 of Gabaix (2009) to deliver price-dividend ratios as a function of the variable portion of resilience.

In the power-utility case above, this was rather straightforward. However, for the EZ case, the tricky part is to deliver M_{t+1} in a "usable" form for the procedure outlined above. The classical Epstein-Zin discount factor has embedded in it the return on a claim to aggregate consumption. In order to use the same intuition presented in this appendix, one can decompose the return on a consumption claim into the resilience of a consumption claim and a i.i.d shock. Gabaix (2012) shows how to do this in both an approximate and complete sense in his appendix. Hence, I omit those derivations from this exposition since the intuition is exactly as before.