

Disaster Risks, Volatility, and Asset Prices

GaOn Kim*

June 2024

Fluctuations in uncertainty are a key driver of asset prices, risk premia, and welfare. This thesis considers two distinct sources of fluctuations in uncertainty: (1) stochastic volatility of shocks occurring at regular frequencies, and (2) stochastic disaster risks pertaining to rare but extreme left tail events. Using a novel calibration strategy relying on VIX and macro data, we quantify the fluctuations in these two sources of uncertainty. We then find that stochastic disaster risks are substantially more important than stochastic volatility for explaining movements in (1) the VIX, (2) the conditional equity risk premia, and (3) welfare fluctuations.

*I am grateful to my supervisor Nicola Borri for his invaluable guidance throughout this master's thesis project, and also to Pierpaolo Benigno for his insightful guidance in the earlier stages of this project. I have also benefitted greatly from the comments of EIEF faculty members during the master's thesis presentation sessions.

1 Introduction

Fluctuations in uncertainty are widely recognized to be a key driver of asset prices, risk premia, and economic welfare. There are abounding examples of sharp fluctuations in uncertainty being forcefully translated into rapid price swings of various assets. However, there are different forms of uncertainty that fluctuate over time, and they can affect asset prices and welfare outcomes in different ways with varying magnitudes. This thesis exploits the information encoded in asset prices to quantify fluctuations in distinct sources of uncertainty, hence contributing to our understanding of the relative importance of different types of uncertainty shocks for fluctuations in overall uncertainty, time-varying risk premia, and economic welfare.

We focus on two key sources of fluctuations in uncertainty: (a) stochastic *volatility* of shocks occurring at regular frequencies, and (b) stochastic *disaster risks* pertaining to rare but extreme left tail events. These two types of uncertainty differ in terms of their magnitudes and frequencies, as the former has smaller magnitudes with higher frequencies, while the latter has large magnitudes but lower frequencies. In addition, they have significantly different implications for asset prices and business cycles. Disaster risk shocks have been shown to be isomorphic to discount factor shocks and generate simultaneous declines in asset prices and aggregate economic activity by discouraging capital investment in RBC models (Gourio, 2012), whereas volatility shocks generate contractions in New Keynesian models through a precautionary savings channel that reduce aggregate demand (Basu and Bundick, 2017). Furthermore, for disaster risk shocks to have significant business cycle effects, it is necessary for the disaster risks to be relevant for both TFP and capital stock depreciation, which Gourio (2012) notes is another contrast with volatility shocks models (Bloom, 2009).

Despite their differences, both types of uncertainty shocks have been widely studied and found to be critical drivers of asset prices, risk premia, welfare, and business cycle fluctuations. The key exercise conducted in this thesis is to horserace these two distinct sources of fluctuations in uncertainty to quantify their relative importance.

A key measure of uncertainty widely used in both academic research and industry is the VIX, which measures the risk-neutral volatility of future S&P 500 stock prices based

on the prices of various put and call options. We will also focus on properties of the VIX to identify the processes for stochastic volatility and stochastic disaster risks. However, the key challenge that we face in conducting our exercise arises from the observation that both elevated disaster risks and volatility can cause an increase in the VIX. Therefore, fluctuations in the VIX alone cannot be used to identify the relative importance of these different types of uncertainty shocks. We overcome this challenge by proposing a calibration strategy based on an analysis that characterizes key differences in how disaster risks shocks and volatility shocks affect the VIX.

We first present an analytical result that characterizes the pass-through of disaster risk shocks and volatility shocks to the VIX. A volatility shock is *directly* translated to the VIX, in the sense that there is no “distortion” in volatility arising when we switch from the physical probability measure to the risk-neutral measure. Conversely, the impact of a disaster risk shock on the VIX is “amplified” under the risk-neutral distribution, so that a small change in disaster risks from an objectively statistical point of view can generate a significant change in the VIX.

This analytical result is then corroborated through a simulation exercise which clearly illustrates the differences in the pass-through of disaster risk shocks versus volatility shocks for the VIX. As suggested by the analytical result, an increase in the VIX caused by an increase in volatility can be detected through an actual increase in volatility from an objectively statistical point of view. Conversely, an increase in the VIX caused by an increase in disaster risks is barely detectable at all from a statistical point of view. Consequently, given a disaster risk shock and a volatility shock that causes the identical response in the VIX, there are stark statistical differences that can be measured objectively.

We then leverage the results above to build a Simulated Method of Moments calibration strategy that sharply identifies the processes for stochastic disaster risks versus stochastic volatility. The centerpiece of our calibration strategy is a simple regression that uses the current level of the VIX to predict the occurrence of future realizations of moderately large consumption shocks. In a world with large volatility shocks relative to disaster risk shocks, the coefficient on this regression should be positive, reflecting the fact that an increase in the VIX is associated with an objectively higher volatility. Conversely, in a world dominated by

disaster risk shocks, fluctuations in the VIX reflect very small changes in disaster risks, leading to a very small coefficient near zero. Therefore, by targeting these regression coefficients in conjunction with other standard moments, we are able to identify the key parameters of interest.

Equipped with the results from this calibration, we then produce quantitative results regarding the relative importance of disaster risk shocks versus volatility shocks for fluctuations in (1) the VIX, (2) conditional equity risk premia, and (3) welfare costs. Our results consistently indicate that disaster risk shocks are quantitatively more important than volatility shocks for all three of these variables, and this is especially the case for fluctuations in the VIX. This provides a quantitative benchmark for understanding the true nature of “uncertainty shocks” as arising more from disaster risk shocks, rather than volatility shocks. In addition, these results raise questions for a common approach to modelling uncertainty shocks with stochastic *volatility* and using the VIX data to estimate the stochastic volatility process (e.g. Basu and Bundick, 2017).

Related Literature. There are two key strands of relevant literature. Firstly, there is a copious literature in option pricing that estimates sophisticated jump-diffusion models for equity prices, incorporating various forms of stochastic volatility and stochastic jump risks. This literature began with the seminal work of Black and Scholes (1973), who provided the famous Black-Scholes formula that relates option prices to their implied volatility under the risk-neutral measure, assuming a constant-volatility lognormal equity price process without jumps. Then, Merton (1976) extended the model to incorporate jumps and Heston (1993) also incorporated stochastic volatility as well. Finally, Duffie et al. (2000) extended the analysis to a very flexible and general class of affine jump-diffusion models that simultaneously allow for stochastic volatility and stochastic jump risks. Building on these theoretical advances, there is a large body of empirical work, such as the seminal work of Broadie et al. (2007), that estimate the processes for stochastic volatility and jump risks over time.

The success of this option pricing literature lies in its ability to closely match option prices across various maturities and over time. They provide successful statistical models for describing the evolution of option and stock prices over time. However, the key limitation

from the perspective of economic interpretation is the lack of meaningful microfoundations. These models are concerned with modelling and estimating the risk-neutral distribution of stock prices directly, which is in contrast to the macro-finance approach of building equilibrium models of the economy with asset prices computed as an equilibrium outcome. This then makes it impossible to understand the objective sources of uncertainty in the economy, as well as the welfare costs arising from such fluctuations.

The second strand of relevant literature is the macro-finance approach to asset pricing, which focuses on *equilibrium* models with optimizing agents and market clearing. This literature began from the seminal work of Lucas (1978), which established the consumption-based framework for asset pricing, and has expanded in various directions as reviewed by Cochrane (2017). In particular, disaster risks was first proposed by Rietz (1988) as a potential resolution of key asset pricing puzzles, including the high equity premium and low risk-free rate (Mehra and Prescott, 1985; Weil, 1989). Empirical evidence for the presence of disaster risks was then provided through an extensive historical data set under Barro (2006) and Barro and Ursúa (2008). Then, Wachter (2013) and Gabaix (2012) extended this to a *stochastic* disaster risks framework to explain, among other asset market phenomena, the high volatility of equity returns. In addition, stochastic volatility – in the sense defined in this thesis – has been widely incorporated throughout this literature as well. For instance, Bansal and Yaron (2004) incorporated stochastic volatility into a long-run risk framework with recursive preferences as in Epstein and Zin (1989), which demonstrated its contribution to the equity premium and excess return volatility. Furthermore, stochastic volatility has been widely incorporated in the macroeconomic literature as well, illustrating that uncertainty shocks can be a source of business cycle fluctuations (Bloom, 2009; Basu and Bundick, 2017; Fernández-Villaverde et al., 2011).

A key point of contrast for the macro-finance approach against the option pricing approach is that the former focuses on much simpler models for reasons of tractability and because structural parameters are much harder to estimate. This has been reflected in the modelling choices of key papers in the literature. For instance, Basu and Bundick (2017) impose only stochastic volatility as the unique source of uncertainty fluctuations in the economy and use VIX and macro data to estimate the process for this stochastic volatility process. As

a separate example, Barro and Liao (2021) estimate stochastic disaster probabilities using options prices, but only allows for fluctuations in disaster probabilities and instead imposes a constant volatility for normal shocks.

Overall, this thesis can be seen as situated in the midway point between these two strands of the literature. It simultaneously incorporates stochastic volatility and stochastic disaster risks in the style of the option pricing literature, but builds an equilibrium model in the macro-finance tradition. Then, we incorporate analytical techniques from the option pricing literature to derive results that will ultimately allow us to identify the key structural parameters of interest – particularly, the relative importance of stochastic volatility versus stochastic disaster risks.

2 Model

This section presents a consumption-based asset pricing model incorporating the two key features of (1) stochastic volatility and (2) stochastic disaster risks. There is a representative agent with Duffie-Epstein (1992) recursive utility who receives endowments and optimally makes consumption and portfolio choices accordingly. In equilibrium, the agent consumes his endowment, which allows us to tractably compute the equilibrium stochastic discount factor and hence price assets. The model is built in continuous time for tractability.

2.1 The Consumption Process

The stochastic process for the consumption stream $\{C_t\}_{t \geq 0}$ is given by the stochastic differential equation below:

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

where C_{t-} denotes the left-hand limit of the consumption at time t . The process B_t is a standard Brownian motion and v_t governs the time-varying volatility of the diffusion part of the consumption process defined above. The stochastic differential equation for volatility v_t is given below:

$$dv_t = \kappa(\bar{v} - v_t) + \sigma_v \sqrt{v_t} dB_{v,t} \quad (2)$$

where $B_{v,t}$ is another standard Brownian motion, independent of B_t .

Furthermore, N_t is a jump process with each jump size equal to $e^{Z_t} - 1$, where Z_t follows a time-invariant distribution denoted by ν . In this model of consumption growth, the realization of a disaster corresponds to a jump in the N_t process. Since disasters are events with very low consumption growth realizations, Z_t is modelled as a negative random variable and will very often have a large magnitude as well. Furthermore, the jump process N_t has time-varying intensity λ_t which follows the stochastic differential equation defined below:

$$d\lambda_t = \kappa(\bar{\lambda} - \lambda_t) + \sigma_\lambda \sqrt{\lambda_t} dB_{\lambda,t} \quad (3)$$

where $B_{\lambda,t}$ is another standard Brownian motion, and we assume that the three Brownian motions ($B_t, B_{v,t}, B_{\lambda,t}$) are mutually independent. We will henceforth refer to v_t as volatility and λ_t as disaster risks at time t .

The stochastic processes for v_t and λ_t above are from Cox et al. (1985) and follow a standard framework for modelling stochastic volatility. The square root term multiplied to the diffusion term ensures that v_t and λ_t do not fall below 0. The stationary distributions for v_t and λ_t as solved by Cox et al. (1985) follows a Gamma distribution with a rightward skew.

2.2 Preferences

The representative agent is assumed to have recursive preferences as in Duffie and Epstein (1992), which is a continuous-time analogue of Epstein and Zin (1989). The general framework of Duffie and Epstein (1992) allows us to freely calibrate risk aversion and the elasticity of intertemporal substitution separately. However, we assume a unit elasticity of intertemporal substitution because it is known to be particularly tractable analytically and is widely studied as a standard benchmark in the literature (Wachter, 2013; Seo and Wachter, 2018).

The lifetime utility of the representative agent is given by

$$V_t = E_t \left[\int_t^\infty f(C_s, V_s) ds \right] \quad (4)$$

$$f(C, V) = \beta(1 - \gamma)V \left[\log C - \frac{1}{1 - \gamma} \log((1 - \gamma)V) \right] \quad (5)$$

where the parameter β is the time discount rate and γ is the risk aversion.

In Appendix A, we derive analytically that the lifetime utility takes the form below:

$$J(C, \lambda, v) = \frac{(\beta^{-1}C)^{1-\gamma}}{1 - \gamma} \exp(a + b\lambda + cv) \quad (6)$$

where C is current consumption level, λ is the current disaster risk, and v is the current volatility. We will show in Appendix A that the current wealth level is proportional to the current level of consumption with a proportionality constant of β^{-1} . Furthermore, the constants (a, b, c) in the expression above are all derived analytically.

Based on the derivations in Appendix A, we can show that lifetime utility is strictly decreasing in λ and v , respectively, given any positive risk aversion γ . Qualitatively, this shows that an increase in uncertainty – whether it is from disaster risks or volatility – leads to a strict decline in welfare. This is a fairly obvious result, but our key interest in the calibration and analysis below will focus on a quantitative evaluation of the importance of shocks in (λ_t, v_t) for welfare fluctuations.

Under Duffie-Epstein preferences specified as in (4), the stochastic discount factor (SDF) is given by

$$\pi_t = \exp \left\{ \int_0^t f_V(C_s, V_s) ds \right\} f_C(C_t, V_t). \quad (7)$$

We can then compute based on (5) and (6) that

$$f_C(C_t, V_t) = \beta^\gamma C_t^{-\gamma} \exp(a + b\lambda_t + cv_t) \quad (8)$$

and

$$f_V(C_t, V_t) = -\beta[b\lambda_t + cv_t + 1] - [(1 - \gamma)\mu + b\kappa_\lambda \bar{\lambda} + c\kappa_v \bar{v}], \quad (9)$$

both of which are now expressed as functions of (C_t, λ_t, v_t) . Thus, based on the above, it is clear that we can apply Ito's lemma to derive the stochastic differential equation for π_t . We

then obtain the result below:

$$\frac{d\pi_t}{\pi_t} = \mu_{\pi,t}dt + \sigma_{\pi,t}^T \begin{pmatrix} dB_t \\ dB_{\lambda,t} \\ dB_{v,t} \end{pmatrix} + (e^{-\gamma Z_t} - 1)dN_t \quad (10)$$

where

$$\mu_{\pi,t} = -\beta - \mu - E_\nu[e^{(1-\gamma)Z} - 1]\lambda_t + \gamma v_t \quad (11)$$

$$\sigma_{\pi,t} = \begin{pmatrix} -\gamma\sqrt{v_t} \\ b\sigma_\lambda\sqrt{\lambda_t} \\ c\sigma_v\sqrt{v_t} \end{pmatrix}. \quad (12)$$

For future reference, it is useful to note how each shock in this model economy translates to movements in the SDF. Firstly, a positive diffusion shock to consumption growth $dB_t > 0$ generates a diffusive decline in π_t , which corresponds to the standard channel in which higher consumption growth generates a decline in marginal utility.

Secondly, a positive shock to disaster risks $dB_{\lambda,t} > 0$ or volatility $dB_{v,t} > 0$ generates an increase in π_t if $\gamma > 1$.¹ This is because whenever $\gamma > 1$, there is a preference for early resolution of uncertainty (because we assume unit elasticity), which in turn implies that negative shocks to lifetime utility directly generates an increase in current marginal utility of wealth, independent of movements in the current level of consumption.

Finally, it is important to notice that a disaster realization generates a very significant upward jump in the marginal utility of wealth (recall that Z_t is a negative random variable, reflecting a large negative shock to consumption). This naturally follows from the idea that a steep decline in consumption growth generates a very large jump in marginal utility.

¹From the derivations in Appendix A, one can clearly see that $b > 0$ whenever $\gamma > 1$.

2.3 Risk-Free Rate

We also suppose that there is a risk-free asset in our economy. As is standard, the risk-free rate r_t^f denotes the (locally deterministic) instantaneous return on holding the risk-free asset over an infinitesimal time interval $[t, t + dt)$. Following standard derivations, one can show that the risk-free rate is given by

$$r_t^f = -E_{t-} \left[\frac{d\pi_t}{\pi_{t-}} \right] = \mu + \beta - \gamma v_t + E_{t-} [e^{-\gamma Z} (e^Z - 1)] \lambda_t. \quad (13)$$

We can observe that the risk-free rate is declining in v_t , which reflects the standard channel of precautionary motives to save. Furthermore, given that Z_t is negative, we can see that the risk-free rate is also declining in λ_t , which once again reflects a precautionary motive to save in the presence of heightened disaster risks that may lead to a spike in the marginal utility of wealth.

2.4 Dividend Strips and Equity Price

An equity claim in this economy is modelled as a claim to a stream of dividends defined by

$$D_t = C_t^\phi \quad (14)$$

for some leverage parameter $\phi \geq 1$.² The case of $\phi = 1$ corresponds to the standard Lucas (1978) tree case of $C_t = D_t$. However, when $\phi > 1$, any shock to consumption is amplified by a factor of ϕ . This reflects the empirical observation that dividends are substantially more volatile than consumption. In particular, it is consistent with the empirical result from Longstaff and Piazzesi (2008) that dividend payments declined substantially more than consumption during the Great Depression, which is a typical example of a disaster realization.

In order to derive an analytical expression for equity prices, it is convenient to begin by considering dividend strips. A period- s dividend strip is defined as an asset that pays out

²This is a standard modelling assumption for equities as used by Gourio (2008) and Wachter (2013) among others.

D_s in period s . Then, the price of this period- s dividend strip at period $t \leq s$ is given by

$$H(D_t, \lambda_t, v_t, s - t) = E_t \left[\frac{\pi_s}{\pi_t} D_s \right]. \quad (15)$$

By extending the solution methodology of Wachter (2013), we can show that the price function H for dividend strips can be found analytically using martingale methods. In particular, we show in Appendix B that

$$H(D_t, \lambda_t, v_t, \tau) = D_t \exp(a_H(\tau) + b_H(\tau)\lambda_t + c_H(\tau)v_t), \quad (16)$$

where (a_H, b_H, c_H) are known in closed form as solutions to Riccati ODEs.

Based on the above dividend strip prices, we can then determine the prices of equities as

$$F_t = E_t \left[\int_t^\infty \frac{\pi_s}{\pi_t} D_s ds \right] = \int_t^\infty H(D_t, \lambda_t, v_t, s - t) ds, \quad (17)$$

where the final equality follows from interchanging the integral and the expectation. Notice that the price-dividend ratio is then given by

$$G(\lambda_t, v_t) = \int_0^\infty \exp(a_H(\tau) + b_H(\tau)\lambda_t + c_H(\tau)v_t) d\tau. \quad (18)$$

Based on the above, we can then use Ito's lemma to fully characterize the evolution of equity prices as

$$\frac{dF_t}{F_t^-} = \mu_{F,t} dt + \sigma_{F,t}^T \begin{pmatrix} dB_t \\ dB_{\lambda,t} \\ dB_{v,t} \end{pmatrix} + (e^{\phi Z_t} - 1) dN_t \quad (19)$$

$$\begin{aligned} \mu_{F,t} = & \phi \left(\mu + \frac{1}{2}(\phi - 1)v_t \right) + \frac{G_{\lambda_t}}{G_t} \kappa(\bar{\lambda} - \lambda_t) + \frac{1}{2} \frac{G_{\lambda\lambda_t}}{G_t} \sigma_\lambda^2 \lambda_t \\ & + \frac{G_{v_t}}{G_t} \kappa(\bar{v} - v_t) + \frac{1}{2} \frac{G_{vv_t}}{G_t} \sigma_v^2 v_t \end{aligned} \quad (20)$$

$$\sigma_{F,t} = \begin{pmatrix} \phi\sqrt{v_t} \\ \frac{G_{\lambda_t}}{G_t}\sigma_\lambda\sqrt{\lambda_t} \\ \frac{G_{v_t}}{G_t}\sigma_v\sqrt{v_t} \end{pmatrix} \quad (21)$$

where G_t is a short-hand notation for $G(\lambda_t, v_t)$ and the partial derivative notations are defined similarly.

We will henceforth impose the following restriction on parameters:

$$\phi \leq 2\gamma. \quad (22)$$

We show in Appendix B that dividend strip prices are weakly decreasing in λ_t and v_t if and only if the inequality above holds. The property that dividend strip prices are declining in disaster risks and volatility is intuitive and clearly reasonable, since the claim to a more uncertain dividend payout should decline in value. Then, given that this property holds, it also follows that the price-dividend ratio is weakly decreasing in λ_t and v_t . Finally, we note that the restriction (22) above is satisfied in virtually all calibrations of models of the type studied in this thesis.

Equipped with the restriction above, we can interpret the impact of various shocks on the equity price in (19). Firstly, it is clear that a positive consumption growth shock $dB_t > 0$ pushes up the equity price. Then, notice that an increase in uncertainty caused by either $dB_{\lambda,t} > 0$ or $dB_{v,t} > 0$ leads to a decline in the equity price. This results from heightened risks for holding equities leading to a decline in their equilibrium prices. Finally, it is clear that a disaster realization will lead to a sharp decline in the equity price. In fact, the disaster impact is amplified by a factor of ϕ which reflects the leverage that dividends have over consumption.

3 Analysis

In this section we present key results that will later be used to provide the theoretical justification for the identification of key parameters in our calibration. The first subsection provides some analytical results for the VIX, and the second subsection presents some simulation results building on the claims proved in the former. All proofs are relegated to Appendix C.

3.1 Analytical Results for the VIX

Our aim is to provide a sharp analytical characterization of the VIX, which will show how the VIX reflects volatility v_t and disaster risks λ_t . Importantly, the result will show that the VIX reflects λ_t in an amplified manner under the risk-neutral measure. Consequently, fluctuations in λ_t that are small from a statistical point of view can generate substantially amplified effects on the VIX.

The raw formula for the VIX represents an aggregation of the prices of put and call options on the S&P500 across various maturities. As shown by Martin (2016), the VIX defined in this way is a measure of the entropy of the aggregate equity price level under the risk-neutral measure. Formally, we may write

$$VIX_t(\tau) = \frac{2}{\tau} L_t^Q \left[\frac{F_{t+\tau}}{F_t} \right] \quad (23)$$

where $L(X) = \log E(X) - E(\log(X))$ represents the entropy of a positive random variable X , and Q is the risk-neutral measure. The entropy of a positive random variable can be broadly considered as a measure of variability of the random variable, since it essentially measures the “Jensen term” under the concave log function. Therefore, one can broadly surmise that both volatility v_t and disaster risks λ_t could play a key role in determining the VIX.

In order to obtain sharper insights regarding the VIX, we must consider the distribution of equity prices under the risk-neutral measure Q . The Radon-Nikodym derivative process

is given by the standard formula

$$E_t \left[\frac{dQ}{dP} \right] = \exp \left(\int_0^t r_s^f ds \right) \frac{\pi_t}{\pi_0} \quad (24)$$

where P is the objective probability measure. Therefore, based on our previous discussion on the SDF π_t , it is clear that the risk-neutral measure assigns greater weight to states of the world with (1) negative consumption growth shocks $dB_t < 0$, (2) positive disaster risk shocks $dB_{\lambda,t} > 0$, (3) positive volatility shocks $dB_{v,t} > 0$, and (4) disaster realizations $dN_t > 0$. In particular, given that the SDF π_t exhibits a large upward jump in the presence of a disaster realization, it is clear that the prospect of disaster risks would be substantially “exaggerated” under the risk-neutral measure.

We now proceed to formally derive a characterization of the stochastic differential equation for equity price F_t under the risk-neutral measure Q . By employing Girsanov’s theorem for semimartingales, we can derive the following lemma:

Lemma 1. The equity price F_t satisfies the following stochastic differential equation:

$$\frac{dF_t}{F_t^-} = \mu_{F,t}^Q dt + \sigma_{F,t}^{Q \text{ T}} \begin{pmatrix} dB_t^Q \\ dB_{\lambda,t}^Q \\ dB_{v,t}^Q \end{pmatrix} + (e^{\phi Z_t} - 1) dN_t \quad (25)$$

where

1. $\mu_{F,t}^Q < \mu_{F,t}$
2. $\sigma_{F,t}^Q = \sigma_{F,t}$
3. $(B_t^Q, B_{\lambda,t}^Q, B_{v,t}^Q)$ constitute a standard Brownian motion
4. Under the Q measure, N_t has jump intensity $\lambda_t^Q = M_\nu(-\gamma)\lambda_t > \lambda_t$

5. Under the \mathbb{Q} measure, the distribution $\nu^{\mathbb{Q}}$ of Z satisfies

$$\frac{d\nu^{\mathbb{Q}}}{d\nu^{\mathbb{P}}} = \frac{e^{-\gamma Z}}{M_{\nu}(-\gamma)}. \quad (26)$$

The results of Lemma 1 are consistent with the intuition that we established above regarding the risk-neutral measure \mathbb{Q} . Firstly, we have suggested above that \mathbb{Q} would overweight the diffusive shocks $dB_t < 0$, $dB_{\lambda,t} > 0$, and $dB_{v,t} > 0$. All of these are shocks that would lead to a decline in the equity price level. Therefore, this is captured in the result of $\mu_{F,t}^{\mathbb{Q}} < \mu_{F,t}$, which is precisely the consequence of overweighting states of the world in which the equity price declines.

Secondly, we noted previously that disaster risks would be severely “exaggerated” under the risk-neutral measure. This is precisely reflected under points 4 and 5 of Lemma 1 above. In fact, there is both an increase in the jump intensity, as well as a distortion of the disaster size distribution that biases it towards more extreme left tail ($Z < 0$) outcomes.

Equipped with the results of Lemma 1, we can then obtain a sharper characterization of the VIX as promised earlier. The key result is presented in the proposition below:

Proposition 1. In the limit of short time horizons, the VIX is equal to the following

$$\lim_{\tau \downarrow 0} VIX_t(\tau) = \|\sigma_{F,t}\|^2 + 2E_{\nu^{\mathbb{Q}}} \left[\sum_{k=2}^{\infty} \frac{(\phi Z)^k}{k!} \right] \lambda_t^{\mathbb{Q}}. \quad (27)$$

This result shows very transparently the components of the VIX. The first term of $\|\sigma_{F,t}\|^2$ is the physical volatility of the diffusion part of the equity price process F_t under the objective probability measure. This part shows up in the VIX without any distortion because as we saw in Lemma 1, the volatility of the diffusion part of the equity price process is unchanged under the risk-neutral measure as $\sigma_{F,t}^{\mathbb{Q}} = \sigma_{F,t}$. Therefore, we can see that the impact of v_t on VIX_t works through a direct channel in which a higher volatility v_t boosts the physical volatility of the equity price process, which in turn is transparently reflected in the VIX measure without any distortion from the risk-neutral measure. It is also notable that $\|\sigma_{F,t}\|^2$ reflects a small component of λ_t as well, which arises from the assumption that the diffusion

component for the stochastic process of λ_t is increasing in the level of λ_t (see (3)).

The second key component of VIX_t is the disaster risks component given by the second term above. It is clear that any increase in λ_t would be substantially amplified by the distortion induced through the risk-neutral measure. Firstly, we know that the intensity λ_t^Q amplifies the true λ_t by a factor of $M_\nu(-\gamma) > 1$ as we saw in Lemma 1. This is quantitatively significant under any standard calibration of disaster risks, and so acts to significantly amplify the effect of λ_t on VIX_t . Furthermore, we know that ν^Q distorts the true disaster size distribution in a way that biases it towards more extreme left-tail outcomes. As a result, we can clearly see that the expectation term would also be distorted upwards under the risk-neutral measure.

Overall, the key takeaway from Proposition 1 above is the following. The impact of v_t on VIX_t works through a direct channel in which there is a higher volatility of equity prices from an objectively statistical point of view. Conversely, the impact of λ_t on VIX_t is severely amplified through a distortion under the risk-neutral measure, so that a small change in λ_t from a statistical point of view can generate a very large increase in the VIX_t . We will later see that these results will be very important for identifying key parameters in our calibration.

3.2 Simulation Results

In this subsection we present some simulation results building upon the analytical results presented in the previous subsection. In particular, the simulation results will provide a quantitative comparison between (1) a volatility shock $v_t \uparrow$, and (2) a disaster risk shock $\lambda_t \uparrow$ which generate an identical response in VIX_t . As suggested by the analytical results, we will show that the volatility shock generates an increase in consumption growth volatility that can be observed and measured from an objective statistical point of view, whereas the disaster risk shock generates an increase in disaster risks which is barely detectable at all, even for long time series samples.

In Appendix D, we combine a log-linearization of the price-dividend ratio with the trans-

form analysis method of Duffie et al. (2000) to show that the VIX takes the form

$$VIX_t^2 = A_0 + A_\lambda \lambda_t + A_v v_t. \quad (28)$$

where $A_0, A_\lambda, A_v > 0$ are constants. Therefore, it follows that λ_t must increase by $1/A_\lambda$ in order to raise VIX_t^2 by one unit, and that v_t must increase by $1/A_v$ to raise VIX_t^2 by one unit.

Based on the above, we propose the following simulation experiment. Let us simulate a long series (1000 months) of log consumption growth based on the following three separate scenarios:

1. the steady state $(\bar{\lambda}, \bar{v})$
2. a high disaster risk state $(\bar{\lambda} + \Delta\lambda, \bar{v})$ with $\Delta\lambda$ chosen so that VIX_t^2 rises by exactly one standard deviation
3. a high volatility state $(\bar{\lambda}, \bar{v} + \Delta v)$ with Δv chosen so that VIX_t^2 rises by exactly one standard deviation

Therefore, the shocks of $\Delta\lambda$ and Δv are comparable in the sense that they induce the identical one-standard-deviation increase in the VIX. However, Figure 1 below shows that the resulting simulated series of log consumption growth are strikingly different.

Figure 1. Disaster Risk Shock vs Volatility Shock for a 1-Std. Increase in VIX




Figure 1 shows simulation results for log consumption growth based on three different situations, defined in the main text. For each situation, we draw 1000 periods of log consumption growth and plot the simulated series in three different panels above. Note that the parameter values used for the simulation are based on the main calibration described under the SMM calibration procedure section below.

It is immediately obvious that the high disaster risk state is essentially indistinguishable from the steady state, whereas the high volatility state is clearly distinguishable from the steady state. This is indeed consistent with the analytical results obtained in the previous section. A slight increase in disaster risks which is very small from a statistical point of view can generate a large increase in the VIX due to the “distortions” embodied in the risk-neutral measure. Conversely, the volatility is translated directly without any distortion into the VIX, so a volatility-induced increase in the VIX must be reflected transparently in the time series of consumption growth over time. These key observations will be the cornerstone for our identification strategy in the next section.

4 Calibration by Simulated Method of Moments

In this section we present the calibration strategy and results. The method of calibration used will be Simulated Method of Moments (SMM). The main topic of interest in this thesis pertains to the quantification of stochastic disaster risks and stochastic volatility. Therefore,

the first subsection is dedicated to presenting the identification strategy for the relative magnitudes of fluctuations in disaster risks versus volatility (i.e. σ_λ vs σ_v). Then, the overall calibration procedure and results will be presented in the subsequent subsection.

Throughout our calibration, we will let 1 unit of time correspond to 1 month in the data.

4.1 Identifying the Relative Magnitudes of σ_λ versus σ_v

We begin by defining the following indicator variable

$$LCS_t = 1 \left(|\Delta \log C_t - \overline{\Delta \log C}| > \text{std}(\widehat{\Delta \log C}) \right) \quad (29)$$

where $\Delta \log C_t$ is the log consumption growth realized at period t ; $\overline{\Delta \log C}$ is the sample mean of log consumption growth; and $\text{std}(\widehat{\Delta \log C})$ is the sample standard deviation of log consumption growth. Essentially, LCS_t is an indicator variable equal to 1 whenever the log consumption growth is more than 1 standard deviation above its mean. This variable is meant to capture *moderately* large consumption growth realizations, not necessarily disasters.

We then consider the following regression:

$$LCS_{t+h} = \alpha_h + \beta_h \log VIX_t + \epsilon_{t+h}. \quad (30)$$

This is a predictive regression which uses the current level of the VIX to predict the likelihood of future large consumption realizations. In particular, our key coefficient of interest will be β_h .

Based on the discussions in the previous section, it is immediately clear that the coefficient β_h will be very sensitive to the relative magnitudes of σ_λ versus σ_v . Suppose, for instance, that $\sigma_\lambda \gg \sigma_v$, which means that we live in a world with stochastic disaster risks but very little fluctuations in volatility. Based on our previous analysis, we know that small changes in disaster risks from a statistical point of view generates large fluctuations in the VIX. Thus, it follows that the coefficient β_h will be close to 0 in this situation. By the converse logic, as σ_v becomes relatively more important than σ_λ , the predictive coefficient β_h will grow larger.

We first begin by running this regression in the data. The regression is based on monthly

data from 1990 January to 2024 January (409 months).³ The results are reported in Panel A of Table 1 below. The coefficient β_h is statistically significant and slowly decays in magnitude as the horizon h becomes longer.

Next, we simulate the model and produce regression results as in Panel B of Table 1 below. The column of main calibration corresponds to the parameters that arise from the main calibration, which is described in the subsection below. This is the result when the parameters are chosen to fit the data, including the regression coefficients shown. However, we also consider two other cases in the next two columns. Firstly, when we halve the value of σ_λ to $\sigma_\lambda/2$ from the main calibration, we find that the coefficients β_h rise substantially, reflecting the fact that stochastic volatility becomes much more important than stochastic disaster risks. The opposite pattern arises when σ_v is halved to $\sigma_v/2$ from the main calibration. Thus, the patterns that we observe over different parameter values are indeed consistent with the theoretical predictions that we discussed above. Overall, the findings from the simulation exercise in Panel B of Table 1 indicate that targeting β_h will sharply identify the relative magnitudes of σ_v versus σ_λ .

³The starting month of 1990 January is the first month in which the CBOE published its VIX data.

Table 1. Predicting Large Consumption Shocks: Sensitivity to σ_λ and σ_v

Panel A: Regression Results (Data)

$$LCS_{t+h} = \alpha_h + \beta_h \log VIX_t + \epsilon_{t+h}$$

	β_h	0.16**	0.13*	0.14**	0.11*	0.08	0.10*
$SE(\beta_h)$		0.08	0.07	0.06	0.06	0.06	0.06

***p<0.01, **p<0.05, *p<0.10

Panel B: Simulated Model Results

$$LCS_{t+h} = \alpha_h + \beta_h \log VIX_t + \epsilon_{t+h}$$

h	Data	Main Calibration	$\sigma_\lambda/2$	$\sigma_v/2$
1	0.16	0.16	0.36	0.06
2	0.13	0.15	0.32	0.06
3	0.14	0.13	0.29	0.05
4	0.11	0.11	0.24	0.04
5	0.08	0.09	0.20	0.03
6	0.10	0.08	0.18	0.03

Table 1 reports the results for the regression of (30). Panel A reports the regression results based on the data. There are 409 months of observations based on monthly data from 1990 January to 2024 January. In addition to β_h estimates, we report Newey-West standard errors as well. Panel B reports the regression results based on simulated data from the model as well. All reported values are β_h coefficients. The Data column is from the data, just as in Panel A, and are included just for comparison. The Main Calibration corresponds to the parameters from the SMM calibration discussed in the next subsection. The $\sigma_\lambda/2$ column halves the value of σ_λ from the main calibration. The $\sigma_v/2$ does the same by halving σ_v from the main calibration. The simulation is based on the mean coefficients drawn from 100 separate simulations of 409 months of data, with each simulation containing 50 burned periods as well.

4.2 SMM and Results

We now describe the SMM procedure used to obtain the results. Firstly, the set of model parameters are divided into two groups. The first group of parameters (see Panel B of Table 2 below) are for those that are previously studied and estimated in the literature, and are simply taken from such previous studies. In particular, the *average* disaster probability $\bar{\lambda}$ is taken from the extensive historical panel study of Barro and Ursua (2008). Also, we impose

for simplicity a normal distribution for the disaster size distribution, and we take the mean and standard deviation parameters (μ_z, σ_z) from Barro and Ursua's (2008) historical study. Also, the risk aversion parameter is set at $\gamma = 4$, which is broadly consistent with the range of values used in the disaster risks asset pricing literature (e.g. Barro, 2006; Wachter, 2013).

Secondly, the set of parameters that are calibrated based on SMM is listed in Panel C of Table 2 below. Each of these parameters are identified by different statistics that are targeted, which are listed under Panels A and B of Table 2. The parameter μ is simply identified by the mean growth rate of consumption, and \bar{v} is identified by the standard deviation of the growth rate of consumption. For the key parameters of $(\sigma_\lambda, \sigma_v)$, the relative magnitudes of these two are identified by targeting $(\beta_h)_{h=1}^6$ as discussed previously. Then, since the standard deviation of the VIX is increasing in both σ_λ and σ_v , these two parameters are jointly identified as well. The mean reversion parameter κ is identified by the first-order autocorrelation of VIX, as well as the rate of decay of $(\beta_h)_{h=1}^6$. Furthermore, the discount rate β is identified by the risk-free rate. Finally, the leverage parameter ϕ is identified by the equity premium and the equity return volatility, both of which are increasing in leverage.

We run the SMM based on the following procedure. First, we draw 100 separate simulations, where each simulation contains a total of 459 periods. Out of the 459 periods, 50 is discarded as part of a burned period, whereas 409 periods are used as simulated data in correspondence with the 409 months of actual data used.⁴ Then, for each of the 100 simulations, the full set of targeted statistics are produced and recorded based on the 409 periods of data. We then compute the average of the produced targeted statistics over the full set of simulations. Finally, we find the parameter values that minimize the squared sum of *percentage* deviations for each of the target statistics. This is taken to be the SMM calibrated parameters reported in Panel D of Table 2 below.

⁴Notice that convergence to the stationary distribution is very fast, as we only require the exogenous stochastic processes for v_t and λ_t to converge to the Gamma stationary distribution.

Table 2. Parametrization

Panel A: Baseline Moments

	Data	Model
$E(\log g)$	1.64	1.59
$\sigma(\log g)$	4.24	4.22
$E(r^f)$	0.31	0.35
$\sigma(r^f)$	0.68	0.88
$E(\log R)$	7.09	7.29
$\sigma(\log R)$	12.60	12.33
$E(\log VIX)$	2.92	3.00
$\sigma(\log VIX)$	0.34	0.35
$AC1(\log VIX)$	0.83	0.88

Panel B: Regression Coefficients

$$LCS_{t+h} = \alpha_h + \beta_h \log VIX_t + \epsilon_{t+h}$$

h	Data	Model
1	0.16	0.16
2	0.13	0.15
3	0.14	0.13
4	0.11	0.11
5	0.08	0.09
6	0.10	0.08

Panel C: Parameters from the Literature

$\bar{\lambda}$	μ_z	σ_z	γ
0.036/12	-0.21	0.13	4

Panel D: Calibrated Parameters from SMM

μ	\bar{v}	σ_v	σ_λ	κ	β	ϕ
0.0013	0.0002	0.0039	0.0248	0.0985	0.0015	2.7827

Table 2 reports the target statistics and parameters for the SMM calibration. Panels A and B report the full set of targeted statistics, both from the actual data as well as the simulated data from the model. All data moments are produced from data for 1990 January to 2023 January (409 months in total). All moments in Panel A are annualized except for those pertaining to the VIX, since the VIX variable by convention is already an annualized quantity. All model moments follow the same conventions for the data. All simulated model statistics are produced from 2000 separate simulations of 409 periods of data, with each simulation including 50 burn periods as well. Panel C reports parameter values taken from the literature, and Panel D reports the calibrated parameters estimated from the SMM procedure. Refer to the main text for additional details on the calibration procedure.

5 Main Results

The main objective of this thesis is to quantify the importance of stochastic volatility and stochastic disaster risks, respectively, for fluctuations in (1) the VIX, (2) conditional equity risk premia, and (3) welfare. Equipped with the calibration results from the previous section, we are now able to present quantitative results pertaining to these questions. The key quantitative results are presented in Table 3 below, with discussions in the following subsections.

Table 3. Main Results: Disaster Risk Shocks vs Volatility Shocks

Panel A: Variance Decomposition for VIX

	All Shocks	Only Disaster Risk Shocks	Only Volatility Shocks
$\text{Var}(VIX_t^2/100)$	12.90	11.97	0.94

Panel B: Variance Decomposition for Conditional Equity Risk Premia

	All Shocks	Only Disaster Risk Shocks	Only Volatility Shocks
$\text{Var}(RP_t)$	16.00	13.86	2.10

Panel C: Consumption-Equivalent Welfare Costs Comparison

C_λ	C_v
1.08%	0.22%

Table 3 summarizes the main results for this thesis. Panel A reports a variance decomposition for the quantity $VIX_t^2/100$. Notice that each period is monthly, and the VIX is annualized just as in the data. The division by 100 is simply a normalization that makes the results quantitatively neater. Also, Panel A is an *exact* decomposition based on an affine decomposition of the VIX. Panel B reports the variance decomposition for the quantity RP_t , which is the annualized risk premium. Notice once again that each period is monthly. Panel B is a non-linear decomposition computed based on a long simulation of RP_t . Finally, Panel C reports the consumption-equivalent welfare costs. C_λ is the consumption-equivalent welfare cost of a 1 (unconditional) standard deviation shock to disaster risks, and C_v is defined similarly. See the main text for a more formal definition.

5.1 Variance Decomposition for the VIX

Based on equation (28) above, we have an affine decomposition of the VIX which can be used to perform an *exact* variance decomposition. The quantitative results are presented in Panel A of Table 3 above. In percentage terms, 92.7% of fluctuations in the VIX are explained by stochastic disaster risks, whereas only 7.3% of fluctuations in the VIX are explained by stochastic volatility.

These results provide a guideline for interpreting the origins of fluctuations in the VIX. In particular, given that most of the fluctuations in the VIX are explained by disaster risk fluctuations, this raises questions about approaches in the literature that interpret VIX fluctuations as arising from stochastic volatility. For instance, Basu and Bundick (2017) build a business cycle model with stochastic volatility in the TFP process and the preference shocks process. Then, they estimate the stochastic volatility process using a combination of VIX and macro data. However, the results in this section suggest that approaches such as this one wrongly assign VIX fluctuations to stochastic volatility, when in fact, they arise primarily from disaster risk fluctuations.

5.2 Variance Decomposition for Conditional Equity Risk Premia

The instantaneous equity return is defined by ⁵

$$dR_t = \frac{dF_t}{F_{t-}} + \frac{D_t}{F_t}. \quad (31)$$

Then, we may define the conditional equity risk premium at time t as

$$RP_t = E_{t-}[dR_t] - r_t^f. \quad (32)$$

It is well-known in the asset pricing literature that the conditional equity risk premium is significantly time-varying, and that fluctuations in uncertainty can be a key driver of this time-variation (Martin, 2016). Consequently, we perform a variance decomposition to examine the relative importance of stochastic disaster risks versus stochastic volatility for

⁵Notice that the price-dividend ratio exhibits no jumps.

explaining this time-varying risk premium.⁶

The results shown in Panel B of Table 2 indicate that disaster risk shocks are once again substantially more important than volatility shocks for explaining time-varying equity risk premia. In fact, disaster risk shocks account for 86.6% of the total variation, whereas volatility shocks account for 13.1%. These results suggest once again that disaster risk shocks are substantially more important for explaining asset market dynamics relative to stochastic volatility.

5.3 Welfare Costs from Disaster Risk vs Volatility Shocks

A consumption-equivalent measure of welfare costs can be developed to quantify the welfare loss from disaster risks versus volatility shocks. In order to define this measure, consider a 1 (unconditional) standard deviation increase in λ_t . Then, we define C_λ as the once-and-for-all increase in log consumption required to exactly compensate for the λ_t increase. Formally, C_λ is implicitly defined as below:

$$J(C \exp(C_\lambda), \lambda + \text{std}(\lambda), v) = J(C, \lambda, v). \quad (33)$$

where $\text{std}(\lambda)$ is the unconditional standard deviation of λ_t . It is notable that the solution for C_λ from above does not depend on the levels of (C, λ, v) . This is a direct consequence of the exponential linear form of the lifetime utility function (see Appendix A).

The welfare cost measure C_v is defined analogously to the above based on a 1 (unconditional) standard deviation shock to v_t . Formally, we have

$$J(C \exp(C_v), \lambda, v + \text{std}(v)) = J(C, \lambda, v). \quad (34)$$

Based on the computation of (C_λ, C_v) reported in Panel C of Table 3, it is clear once again that disaster risk shocks induce a greater loss of welfare than volatility shocks. The difference is substantial, with the welfare cost being about 5 times higher for disaster risk

⁶Unlike the affine equation for the VIX above, we have a non-linear variance decomposition in the case of the conditional risk premium. Therefore, as seen in Panel B of Table 2, the decomposed variance from each component does not add up to the total variance, although it is very close to being exact.

shocks relative to volatility shocks. Therefore, the results from this exercise suggest that uncertainty-induced welfare fluctuations arise more due to disaster risk fluctuations than volatility shocks. However, it is also notable that the extent to which disaster risk shocks are more important than volatility shocks is actually much lower than the relative magnitudes absorbed in the variance decomposition exercises above.

6 Conclusion

Fluctuations in uncertainty are known to be a key driver of asset prices, risk premia, and welfare costs. However, there are several different forms of uncertainty, and the relative importance of these different types of uncertainty shocks are not clear. This thesis has focused on two particular forms of uncertainty shocks – stochastic volatility and stochastic disaster risks – which differ along the dimensions of magnitude and frequency. The overall aim of this thesis was to horserace these two sources of fluctuations in uncertainty in order to examine their relative importance.

In order to identify the relative magnitudes of volatility vs disaster risk shocks, we established key results characterizing the differential pass-through of these two shocks for the VIX. The key result was that an increase in the VIX caused by a volatility shock will exhibit a corresponding increase in objective volatility as measured statistically, whereas the same increase in the VIX caused by a disaster risk shock would correspond to only a small change in disaster risks when measured from a statistical point of view. Therefore, by targeting the coefficients from a predictive regression of the current level of the VIX on future volatility of consumption growth, we are able to sharply identify the relative magnitudes of volatility vs disaster risk shocks.

The key quantitative conclusion which then arises is that stochastic disaster risks are substantially more important than stochastic volatility for explaining variations in (1) the VIX, (2) time-varying conditional equity risk premia, and (3) welfare costs. This is especially the case for explaining fluctuations in the VIX. The results provided in this thesis then provide a quantitative benchmark for understanding the true nature of “uncertainty shocks” as arising from disaster risk fluctuations rather than volatility fluctuations. Furthermore,

it raises questions about approaches in which VIX data is used to estimate the stochastic volatility process without accounting for disaster risk fluctuations.

References

- Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509.
- Barro, R. J. (2006). Rare Disasters and Asset Markets in the Twentieth Century*. *The Quarterly Journal of Economics*, 121(3):823–866.
- Barro, R. J. and Liao, G. Y. (2021). Rare disaster probability and options pricing. *Journal of Financial Economics*, 139(3):750–769.
- Barro, R. J. and Ursúa, J. F. (2008). Consumption disasters in the twentieth century. *American Economic Review*, 98(2):58–63.
- Basu, S. and Bundick, B. (2017). Uncertainty shocks in a model of effective demand. *Econometrica*, 85(3):937–958.
- Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.
- Broadie, M., Chernov, M., and Johannes, M. (2007). Model specification and risk premia: Evidence from futures options. *The Journal of Finance*, 62(3):1453–1490.
- Cochrane, J. H. (2017). Macro-Finance*. *Review of Finance*, 21(3):945–985.
- Cox, J. C., Ingersoll, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407.
- Duffie, D. and Epstein, L. G. (1992). Stochastic differential utility. *Econometrica*, 60(2):353–394.
- Duffie, D., Pan, J., and Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376.
- Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969.

- Fernández-Villaverde, J., Guerrón-Quintana, P., Rubio-Ramírez, J. F., and Uribe, M. (2011). Risk matters: The real effects of volatility shocks. *The American Economic Review*, 101(6):2530–2561.
- Gabaix, X. (2012). Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance *. *The Quarterly Journal of Economics*, 127(2):645–700.
- Gourio, F. (2012). Disaster risk and business cycles. *American Economic Review*, 102(6):2734–66.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327–343.
- Lucas, R. E. (1978). Asset prices in an exchange economy. *Econometrica*, 46(6):1429–1445.
- Martin, I. (2016). What is the Expected Return on the Market?*. *The Quarterly Journal of Economics*, 132(1):367–433.
- Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161.
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1):125–144.
- Rietz, T. A. (1988). The equity risk premium a solution. *Journal of Monetary Economics*, 22(1):117–131.
- Seo, S. B. and Wachter, J. A. (2018). Do rare events explain cdx tranche spreads? *The Journal of Finance*, 73(5):2343–2383.
- Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, 68(3):987–1035.
- Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24(3):401–421.

Appendix

A Lifetime Utility Function

The derivation of the lifetime utility function is based on the approach of Wachter (2013), which has been extended to fit our current model context.

The representative agent in the model is entitled to the consumption endowment process given by (1) in the main text. Let us consider a “consumption claim” asset that provides a stream of payments equal to the stochastic process for $\{C_t\}_{t \geq 0}$. Then, denote the price of this consumption claim as S_t . In equilibrium, the price S_t will be equal to the total wealth of the representative agent at any point in time. It will be convenient to denote the wealth separately as W_t .

Let us then define $V(W_t, \lambda_t, v_t)$ as the value function (i.e. lifetime utility function) of the agent given current wealth W_t , current disaster risk λ_t , and current volatility v_t . The reason we work with the wealth W_t rather than C_t is because “standard” functional forms for guessing the lifetime utility function in these contexts is based on the wealth, rather than current consumption level.

Then, we will let the agent face a portfolio choice problem between two assets: (1) the consumption claim asset defined above, with price S_t , and (2) the risk-free asset with instantaneous rate r_t . In equilibrium the agent will simply hold 1 unit of the consumption claim and hold 0 units of the risk-free asset. Therefore, the availability of the risk-free asset does not affect the agent’s lifetime utility function in any way – however, the reason we do include it here is because the optimal portfolio choice condition between the risk-free asset and the consumption claim actually helps us derive the proportion between consumption and wealth, which is useful for deriving the lifetime utility function in terms of consumption (this is necessary to derive the SDF in the main text). In addition, it is worth noting that although we do not model the portfolio choice problem with respect to other assets (e.g. equities, options, and long-term bonds), this does not mean that we are unable to price these assets within our framework. Equipped with the equilibrium SDF derived in the main text, we can price any well-defined asset.

We will begin with the conjecture, later verified, that the wealth-consumption ratio is a constant, denoted l below:

$$\frac{W_t}{C_t} = l. \quad (35)$$

This guess arises from previous results (e.g. Weil, 1989) which show that representative agent models with recursive preferences and $EIS = 1$ tend to have a constant wealth-consumption ratio.

Given the conjecture above, we know then that the stochastic process for S_t must satisfy the following in equilibrium:

$$\frac{dS_t}{S_t} = \mu dt + \sqrt{v_t} dB_t + (e^{Z_t} - 1) dN_t. \quad (36)$$

Then, denoting the fraction of wealth invested in the consumption claim as α_t , it follows that wealth evolves according to

$$dW_t = [\alpha_t(\mu + l^{-1}) + (1 - \alpha_t)r_t]W_t dt - C_t dt + W_t \alpha_t \sqrt{v_t} dB_t + W_t \alpha_t (e^{Z_t} - 1) dN_t. \quad (37)$$

We are then able to formulate the HJB equation for the value function as below:

$$\begin{aligned} 0 = & \sup_{C_t, \alpha_t} f(C_t, V_t) + V_{W_t} [W_t(\alpha_t(\mu + l^{-1}) + (1 - \alpha_t)r_t) - C_t] + \frac{1}{2} V_{WW_t} W_t^2 \alpha_t^2 v_t \\ & + \lambda_t E_\nu [V(W_t(1 + \alpha_t(e^{Z_t} - 1)), \lambda_t, v_t) - V(W_t, \lambda_t, v_t)] + V_{\lambda_t} \kappa_\lambda (\bar{\lambda} - \lambda_t) \\ & + \frac{1}{2} V_{\lambda\lambda_t} \sigma_\lambda^2 \lambda_t + V_{v_t} \kappa_v (\bar{v} - v_t) + \frac{1}{2} V_{vv_t} \sigma_v^2 v_t \end{aligned} \quad (38)$$

where V_t is used as short-hand notation for $V(W_t, \lambda_t, v_t)$; V_{W_t} is used as short-hand notation for $V_W(W_t, \lambda_t, v_t)$; and similarly for other partial derivatives with a time index.

Let us then substitute the equilibrium conditions $\alpha_t = 1$ and $C_t = l^{-1}W_t$. Then, we know that the following HJB equation must hold:

$$\begin{aligned} 0 = & f(l^{-1}W_t, V_t) + V_{W_t} W_t \mu + \frac{1}{2} V_{WW_t} W_t^2 v_t \\ & + \lambda_t E_\nu [V(W_t e^{Z_t}, \lambda_t, v_t) - V(W_t, \lambda_t, v_t)] + V_{\lambda_t} \kappa_\lambda (\bar{\lambda} - \lambda_t) \\ & + \frac{1}{2} V_{\lambda\lambda_t} \sigma_\lambda^2 \lambda_t + V_{v_t} \kappa_v (\bar{v} - v_t) + \frac{1}{2} V_{vv_t} \sigma_v^2 v_t. \end{aligned} \quad (39)$$

Let us then guess the functional form below, which is a common guess for $EIS = 1$:

$$V(W_t, \lambda_t, v_t) = \frac{W_t^{1-\gamma}}{1-\gamma} I(\lambda_t, v_t). \quad (40)$$

Now let us consider the FOC with respect to consumption in the HJB (38). Upon algebraic manipulation, we can check that this FOC holds true in equilibrium if and only if

$$\beta = l^{-1}. \quad (41)$$

Now, let us plug in the guess (40) together with (41) into the HJB of (39). Upon algebraic manipulation, we find that the guess is verified if and only if the following ODE for $I(\lambda_t, v_t)$ is satisfied:

$$\begin{aligned} \log I(\lambda_t, v_t) = & (1-\gamma) \log \beta + \frac{1}{\beta} (1-\gamma) \mu - \frac{1}{2} \frac{1}{\beta} \gamma (1-\gamma) v_t + \frac{\lambda_t}{\beta} E_\nu [e^{(1-\gamma)Z} - 1] \\ & + \frac{1}{\beta} \frac{I_{\lambda_t}}{I_t} \kappa_\lambda (\bar{\lambda} - \lambda_t) + \frac{1}{2} \frac{1}{\beta} \frac{I_{\lambda\lambda_t}}{I_t} \sigma_\lambda^2 \lambda_t \\ & + \frac{1}{\beta} \frac{I_{v_t}}{I_t} \kappa_v (\bar{v} - v_t) + \frac{1}{2} \frac{1}{\beta} \frac{I_{vv_t}}{I_t} \sigma_v^2 v_t \end{aligned} \quad (42)$$

where I_t is used as short-hand notation for $I(\lambda_t, v_t)$ and similarly for the partial derivatives with the time index.

Then, let us guess the exponential linear functional form as the solution to the ODE above:

$$I(\lambda_t, v_t) = \exp(a + b\lambda_t + cv_t). \quad (43)$$

Plugging the above into (42), we find that the guess is verified if and only if the following hold:

$$a = (1-\gamma) \log \beta + \frac{1}{\beta} (1-\gamma) \mu + \frac{1}{\beta} b \kappa_\lambda \bar{\lambda} + \frac{1}{\beta} c \kappa_v \bar{v} \quad (44)$$

$$b = \frac{\kappa_\lambda + \beta \pm \sqrt{(\kappa_\lambda + \beta)^2 - 2\sigma_\lambda^2 E_\nu [e^{(1-\gamma)Z} - 1]}}{\sigma_\lambda^2}. \quad (45)$$

$$c = \frac{\kappa_v + \beta \pm \sqrt{(\kappa_v + \beta)^2 - \sigma_v^2 \gamma (1-\gamma)}}{\sigma_v^2}. \quad (46)$$

Notice that we have two candidates for b and c , respectively. However, we can easily rule out

the + solution based on the following economic reasoning. Suppose that the disaster size is always equal to 0 with probability 1 (i.e. $Z \equiv 0$). In that case, disaster risk fluctuations should have no impact at all on the lifetime utility function (i.e. $b = 0$). The solution of b consistent with this is the – solution. In the case of c , consider the limiting case as $\gamma \rightarrow 0$ so that the agent approaches risk-neutrality. It follows that volatility risk fluctuations should have no impact at all on agent’s lifetime utility (notice that v_t has no impact at all on future expected consumption). The only solution consistent with this property is the case with the – solution for c . Thus, we are able to pin down the unique solution for b and c .

Now, the only optimality that must be verified to complete our verification is the FOC with respect to the risk-free rate. It is clear that this is indeed verified if and only if

$$r_t = \mu + \beta - \gamma v_t + \lambda_t E_\nu[e^{-\gamma Z}(e^Z - 1)]. \quad (47)$$

Notice further that this risk-free rate is indeed consistent with the risk-free rate in the main text which was derived based on the SDF.

This completes our derivation of the lifetime utility function $V(W, \lambda, v)$. However, it will be convenient for us to express the lifetime utility function in terms of current consumption level C_t when deriving the SDF. Based on (41), we can easily express with lifetime utility function as

$$J(C, \lambda, v) = \frac{(\beta^{-1}C)^{1-\gamma}}{1-\gamma} \exp(a + b\lambda + cv). \quad (48)$$

B Dividend Strips

In this section, we adapt the methodology of Wachter (2013) to our model, presenting an analytical solution for the dividend strip price function. Firstly, let us fix some $s > 0$ and view $\{H(D_t, \lambda_t, v_t, s - t)\}_{0 \leq t \leq s}$ as a stochastic process with time index t . At times we will find it convenient to denote this process as H_t . Then, notice immediately that

$$\pi_t H(D_t, \lambda_t, v_t, s - t) = E_t[\pi_s H(D_s, \lambda_s, v_s, 0)], \quad (49)$$

from which it follows that $\{\pi_t H(D_t, \lambda_t, v_t, s - t)\}_{0 \leq t \leq s}$ – viewed as a stochastic process with time index t – is a martingale.

Let us then conjecture the following functional form for the dividend strips prices:

$$H(D_t, \lambda_t, v_t, \tau) = D_t \exp\{a_H(\tau) + b_H(\tau)\lambda_t + c_H(\tau)v_t\}. \quad (50)$$

Then, we may apply Ito's lemma to characterize the stochastic differential equation for H_t . We then obtain the following:

$$\frac{dH_t}{H_{t-}} = \mu_{H,t} dt + \sigma_{H,t}^T \begin{pmatrix} dB_t \\ dB_{\lambda,t} \\ dB_{v,t} \end{pmatrix} + (e^{\phi Z_t} - 1) dN_t \quad (51)$$

$$\begin{aligned} \mu_{H,t} = & - [a'_H(s-t) + b'_H(s-t)\lambda_t + c'_H(s-t)v_t] + \mu_{D,t} \\ & + b_H(s-t)\kappa_\lambda(\bar{\lambda} - \lambda_t) + \frac{1}{2}b_H^2(s-t)\sigma_\lambda^2\lambda_t \\ & + c_H(s-t)\kappa_v(\bar{v} - v_t) + \frac{1}{2}c_H^2(s-t)\sigma_v^2v_t \end{aligned} \quad (52)$$

$$\mu_{D,t} = \phi\left(\mu + \frac{1}{2}(\phi - 1)v_t\right) \quad (53)$$

$$\sigma_{H,t} = \begin{pmatrix} \phi\sqrt{v_t} \\ b_H(s-t)\sigma_\lambda\sqrt{\lambda_t} \\ c_H(s-t)\sigma_v\sqrt{v_t} \end{pmatrix}. \quad (54)$$

Then, consider the stochastic process for $\{\pi_t H_t\}_{0 \leq t \leq s}$. We may apply Ito's lemma to derive the stochastic differential equation for this process. Then, since we know that this process must be a martingale, we can impose the following condition:

$$E_{t-}[d(\pi_t H_t)] = 0, \quad (55)$$

which in turn translates to the condition

$$\mu_{\pi,t} + \mu_{H,t} + \sigma_{\pi,t}^T \sigma_{H,t} + E_\nu[e^{(\phi-\gamma)Z_t} - 1]\lambda_t = 0. \quad (56)$$

We can express the left hand side as an affine function of (λ_t, v_t) . Thus, the condition above corresponds to three separate conditions, in which we set the constant component and the two coefficients on (λ_t, v_t) as equal to 0. These three respective conditions are listed below:

$$a'_H(\tau) = -\beta - \mu + \phi\mu + b_H(\tau)\kappa_\lambda\bar{\lambda} + c_H(\tau)\kappa_v\bar{v} \quad (57)$$

$$b'_H(\tau) = \frac{1}{2}\sigma_\lambda^2 b_H^2(\tau) + (b\sigma_\lambda^2 - \kappa_\lambda)b_H(\tau) + E_\nu[e^{(\phi-\gamma)Z} - e^{(1-\gamma)Z}] \quad (58)$$

$$c'_H(\tau) = \frac{1}{2}\sigma_v^2 c_H^2(\tau) + [c\sigma_v^2 - \kappa_v]c_H(\tau) + \left(\frac{1}{2}\phi - \gamma\right)(\phi - 1). \quad (59)$$

The three equations constitute a system of Riccati ODEs with known analytical solutions. The solutions are presented below:

$$b_H(\tau) = 2E_\nu[e^{(\phi-\gamma)Z} - e^{(1-\gamma)Z}] \frac{e^{-\Delta_{b_H}\tau} - 1}{(b\sigma_\lambda^2 - \kappa + \Delta_{b_H})(1 - e^{-\Delta_{b_H}\tau}) - 2\Delta_{b_H}} \quad (60)$$

$$\Delta_{b_H} = \sqrt{(b\sigma_\lambda^2 - \kappa)^2 - 2\sigma_\lambda^2 E_\nu[e^{(\phi-\gamma)Z} - e^{(1-\gamma)Z}]} \quad (61)$$

$$c_H(\tau) = (\phi - 2\gamma)(\phi - 1) \frac{e^{-\Delta_{c_H}\tau} - 1}{(c\sigma_v^2 - \kappa + \Delta_{c_H})(1 - e^{-\Delta_{c_H}\tau}) - 2\Delta_{c_H}} \quad (62)$$

$$\Delta_{c_H} = \sqrt{(c\sigma_v^2 - \kappa)^2 - \sigma_v^2(\phi - 2\gamma)(\phi - 1)} \quad (63)$$

$$\begin{aligned} a_H(\tau) = & (-\beta - \mu + \phi\mu)\tau \\ & + 2E_\nu[e^{(\phi-\gamma)Z} - e^{(1-\gamma)Z}]\kappa_\lambda\bar{\lambda}\Gamma(\tau; \Delta_{b_H}, b\sigma_\lambda^2 - \kappa) \\ & + (\phi - 2\gamma)(\phi - 1)\kappa_v\bar{v}\Gamma(\tau; \Delta_{c_H}, c\sigma_v^2 - \kappa) \end{aligned} \quad (64)$$

where the Γ is a function defined as below:

$$\Gamma(t; \Delta, \alpha) = \frac{2}{\Delta^2 - \alpha^2} \log \left(\left[\frac{\alpha}{2\Delta} + \frac{1}{2} \right] e^{-\Delta t} + \frac{1}{2} - \frac{\alpha}{2\Delta} \right) + \frac{t}{\Delta - \alpha}. \quad (65)$$

We conclude this section by revisiting the claim in Section 2 that dividend strip prices

are weakly decreasing in λ_t and v_t if and only if $\phi \leq 2\gamma$. To see this result, it is sufficient to note that dividend strip prices are decreasing in λ_t if and only if $b_H(\tau) \leq 0$ for all $\tau \geq 0$; and dividend strip prices are decreasing in v_t if and only if $c_H(\tau) \leq 0$ for all $\tau \geq 0$. In fact, we can see that since $\phi \leq 1$ and $Z < 0$ with probability 1, the former condition always holds. In addition, we can show upon algebraic manipulations that $c_H(\tau) \leq 0$ for all $\tau \geq 0$ if and only if $\phi \leq 2\gamma$ holds. This establishes the claim from the main text.

C Proofs

Proof of Lemma 1. By applying Ito's lemma, we can find that

$$d \log \pi_t = \left(\mu_{\pi,t} - \frac{1}{2} \|\sigma_{\pi,t}\|^2 \right) dt + \sigma_{\pi,t}^T \begin{pmatrix} dB_t \\ dB_{\lambda,t} \\ dB_{v,t} \end{pmatrix} - \gamma Z_t dN_t, \quad (66)$$

from which it follows that

$$\frac{\pi_t}{\pi_0} = \exp \left(\int_0^t \left(\mu_{\pi,s} - \frac{1}{2} \|\sigma_{\pi,s}\|^2 \right) ds + \int_0^t \sigma_{\pi,s}^T \begin{pmatrix} dB_t \\ dB_{\lambda,t} \\ dB_{v,t} \end{pmatrix} + \int_0^t \int_{\mathbb{R}_-} (-\gamma Z) \theta(ds \times dz) \right), \quad (67)$$

where θ is the random counting measure associated with $Z_t dN_t$.

Furthermore, we may write the risk-free rate as

$$r_t^f = -E_t \left[\frac{d\pi_t}{\pi_t} \right] = -\mu_{\pi,t} - \lambda_t E_\nu [e^{-\gamma Z} - 1]. \quad (68)$$

Then, recall that the Radon-Nikodym derivative process for the risk-neutral measure

with respect to the physical probability measure can be written as

$$\begin{aligned}
L_t &= \exp\left(\int_0^t r_s^f ds\right) \frac{\pi_t}{\pi_0} \\
&= \exp\left(-\frac{1}{2}\int_0^t \|\sigma_{\pi,s}\|^2 ds + \int_0^t \sigma_{\pi,s}^T \begin{pmatrix} dB_t \\ dB_{\lambda,t} \\ dB_{v,t} \end{pmatrix}\right) \\
&\quad + \int_0^t \int_{\mathbb{R}_-} (-\gamma Z)\theta(ds \times dz) + \int_0^t \int_{\mathbb{R}_-} (1 - e^{-\gamma Z})\iota_s(dz) ds,
\end{aligned} \tag{69}$$

where ι is the intensity kernel associated with θ .

Based on the formulation above, we are ready to apply Girsanov's theorem for semi-martingales. This immediately allows us to obtain the following results:

$$\begin{pmatrix} B_t^Q \\ B_{\lambda,t}^Q \\ B_{v,t}^Q \end{pmatrix} \equiv \begin{pmatrix} B_t \\ B_{\lambda,t} \\ B_{v,t} \end{pmatrix} - \int_0^t \sigma_{\pi,s} ds \tag{70}$$

is a standard Brownian motion under \mathbb{Q} . In addition, if we define ι^Q to be the \mathbb{Q} -intensity kernel, we have the following:

$$\iota^Q(\mathbb{R}) = E_\nu[e^{-\gamma Z}]\lambda_t \tag{71}$$

$$\frac{\iota^Q(dz)}{\iota^{\mathbb{R}}} = \frac{e^{-\gamma Z}}{E_\nu[e^{-\gamma Z}]}\nu(dz). \tag{72}$$

Based on the results above, the claims in 2-5 are immediately obvious. Then, we can observe that

$$\mu_{F,t}^Q = \mu_{F,t} + \sigma_{F,t}^T \sigma_{\pi,t}. \tag{73}$$

Finally, we can easily verify that $\sigma_{F,t}^T \sigma_{\pi,t} < 0$ so that the claim in 1 holds as well. \square

Proof of Proposition 1. By an application of Ito's Lemma, we have that

$$d \log F_t = (\mu_{F,t}^Q - \frac{1}{2} \|\sigma_{F,t}\|^2) dt + \sigma_{F,t} \begin{pmatrix} dB_t^Q \\ dB_{\lambda,t}^Q \\ dB_{v,t}^Q \end{pmatrix} + \phi Z_t dN_t. \quad (74)$$

Then, let us consider the limit

$$\begin{aligned} \lim_{\tau \downarrow 0} VIX_t^2(\tau) &= 2 \lim_{\tau \downarrow 0} \frac{L_t^Q \left(\frac{F_{t+\tau}}{F_t} \right)}{\tau} \\ &= \|\sigma_{F,t}\|^2 + 2 \lim_{\tau \downarrow 0} \frac{\log E_t^Q \exp(\phi Z(N_{t+\tau} - N_t)) - E_t^Q \phi Z(N_{t+\tau} - N_t)}{\tau} \\ &= \|\sigma_{F,t}\|^2 + 2E_{\nu^Q} \left[e^{\phi Z} - (1 + \phi Z) \right] \lambda_t^Q, \end{aligned} \quad (75)$$

which is then clearly equal to the expression in the proposition given a Maclaurin series expansion of the exponential. \square

D Computing the VIX

We begin with a log-linearization of the price-dividend ratio around the steady state $(\bar{\lambda}, \bar{v})$:

$$\begin{aligned} \log G(\lambda_t, v_t) &= \log G(\bar{\lambda}, \bar{v}) + \frac{G_\lambda(\bar{\lambda}, \bar{v})}{G(\bar{\lambda}, \bar{v})} (\lambda_t - \bar{\lambda}) + \frac{G_v(\bar{\lambda}, \bar{v})}{G(\bar{\lambda}, \bar{v})} (v_t - \bar{v}) \\ &= a_g + b_g (\lambda_t - \bar{\lambda}) + c_g (v_t - \bar{v}), \end{aligned} \quad (76)$$

where (a_g, b_g, c_g) are constants defined implicitly above.

The key advantage of this log-linearization is that it allows us to cast the evolution of stock prices F_t into the affine jump-diffusion framework of Duffie et al. (2000), under which there are powerful computational and analytical tools. Notice that under this log-linearization, we

may write

$$\frac{F_{t+\tau}}{F_t} = \frac{G_{t+\tau}}{G_t} \left(\frac{C_{t+\tau}}{C_t} \right)^\phi = \exp \left(b_g(\lambda_{t+\tau} - \lambda_t) + c_g(v_{t+\tau} - v_t) + \phi(\log C_{t+\tau} - \log C_t) \right). \quad (77)$$

Then, notice that the X_τ defined below satisfies the assumption of affine jump-diffusion as defined under Duffie:

$$X_\tau = \begin{pmatrix} \log C_{t+\tau} - \log C_t \\ \lambda_{t+\tau} \\ v_{t+\tau} \end{pmatrix}. \quad (78)$$

Finally, notice that the expectations $E_t^Q \left(\frac{F_{t+\tau}}{F_t} \right)$ and $E_t^Q \left[\log \frac{F_{t+\tau}}{F_t} \right]$ can be computed tractably using Duffie et al.'s (2000) transform analysis tools. Therefore, the risk-neutral entropy can be readily calculated, and it is easy to show using Duffie et al.'s (2000) results that the VIX takes the form below

$$VIX_t^2 = A_0 + A_\lambda \lambda_t + A_v v_t. \quad (79)$$