

Option prices in a model with stochastic disaster risk *

Sang Byung Seo

Jessica A. Wachter

University of Houston

University of Pennsylvania

and NBER

October 31, 2016

Abstract

Contrary to well-known asset pricing models, volatilities implied by equity index options exceed realized stock market volatility and exhibit a pattern known as the volatility skew. We explain both facts using a model that can also account for the mean and volatility of equity returns. Our model assumes a small risk of economic disaster that is calibrated based on international data on large consumption declines. We allow the disaster probability to be stochastic, which turns out to be crucial to the model's ability both to match equity volatility and to reconcile option prices with macroeconomic data on disasters.

*First draft: January 18, 2013. Seo: sseo@bauer.uh.edu; Wachter: jwachter@wharton.upenn.edu. We thank Robert Barro, Hui Chen, Domenico Cuoco, Mikhail Chernov, Xavier Gabaix, Ivan Shaliastovich, Viktor Todorov and seminar participants at the NBER Summer Institute, the Johns Hopkins and AQR Conference on the Role of Options in Asset Pricing, Carnegie Mellon University, Cornell University and the Wharton school for helpful comments.

1 Introduction

Prices of equity index options offer a powerful means of discriminating among theories of asset prices. Since the work of Mehra and Prescott (1985) and Campbell and Shiller (1988), asset pricing research has sought to reconcile the apparent low volatility of U.S. consumption growth with the high equity premium and the high volatility of stock prices. One classic theory, presented by Campbell and Cochrane (1999), is that small variation in consumption is very costly to investors, who evaluate current consumption relative to the recent past. Bansal and Yaron (2004) present a second classic theory: while shocks to realized consumption may be small, the risk in the conditional mean and volatility of the consumption distribution is considerable. These theories are highly successful in explaining the mean and volatility of stock returns. Yet option prices remain a challenge.

We illustrate the challenge with the classic graph of option-implied volatilities as a function of the exercise price. Figure 1 shows average implied volatilities for 3-month put options written on the S&P 500 over the 1996 to 2012 period.¹ On the x -axis is the scaled exercise price (the exercise price divided by the index price, known as “moneyness”). Figure 1 also shows the level of annual stock price volatility over this sample, 16%, and the value of annual consumption volatility, below 2%.² Thus, while the volatility of stock prices is high, implied volatilities are even higher, particularly for out-of-the-money put options (those with relatively low exercise prices). This fact that implied volatilities are decreasing in the exercise price is known as the *implied volatility skew* (or, if there is no ambiguity, as the volatility skew), and has been known to the literature since Rubinstein (1994) and Coval and Shumway

¹The implied volatility is defined to be the value that correctly prices the option under the Black and Scholes (1973) model. Implied volatilities are a convenient normalization of option prices that illustrate deviations from lognormality, because, under a lognormal distribution, they would be constant as a function of maturity and exercise price. However, matching implied volatilities is equivalent to matching option prices.

²These are measured as the standard deviation of the log of year-over-year stock prices or year-over-year consumption.

(2001). Figure 1 illustrates three hurdles an asset pricing model must clear to explain these data. First, the model must generate sufficient volatility of asset prices, given low levels of consumption risk. Second, the model must generate even higher implied volatilities, and third, the model must generate yet higher volatilities for options that are further out of the money.

The models of Campbell and Cochrane (1999) and Bansal and Yaron (2004) are able to meet the first of these challenges. However, they cannot meet the second or third for the simple reason that these models are conditionally lognormal over a non-infinitesimal interval. By definition, the conditions of Black and Scholes (1973) hold over the decision interval of economic agents in these models. If this is one quarter, then three-month implied volatility curves must be flat and at the level of realized price volatility. One could extend these models to continuous time; however, unless one fundamentally changed the calibration, this would not generate sufficient non-lognormality to explain the level or slope of the implied volatility curve.³

A model that seems, at first glance, to be consistent with the implied volatility skew, is one that explicitly takes into account conditional non-lognormality in underlying consumption and hence in stock prices. Barro (2006) shows that a model that allows for rare economic disasters can explain the equity premium puzzle of Mehra and Prescott (1985). Because put options pay off in disaster states, they should trade at a premium relative to their value under conditional lognormality, and the premium should be especially large for out-of-the-money puts. We would expect to see an implied volatility skew in such a model. Backus, Chernov, and Martin (2011) show that, when calibrated to macroeconomic data, this model does indeed imply a volatility skew. However, the skew is qualitatively different than in

³This is not to suggest that these models cannot be extended to account for options data. Doing so, however, requires a mechanism in addition to those already in the models that account for the equity premium and volatility (Benzoni, Collin-Dufresne, and Goldstein, 2011; Drechsler and Yaron, 2011; Du, 2011). See Section 4.9 for further discussion.

the data; most notably the slope of the curve is far steeper in the model than in the data counterpart. Their findings appear to demonstrate a basic inconsistency between disasters as reflected in macroeconomic data and non-lognormality as reflected in options markets. In doing so, they call into question the plausibility of disaster risk as an explanation of the equity premium puzzle.

In this paper, we present a rare-disaster model that explains the implied volatility skew. In so doing, we reconcile option prices with the disaster distribution in macroeconomic data. We make one simple change to the rare disaster framework of Barro (2006) and Backus, Chernov, and Martin (2011): we allow the probability of a disaster to vary over time. This change, besides allowing the model to explain option prices, also allows the model to explain realized stock market volatility in normal times. Importantly, the exact same change allows the model to explain both features of the data. Iid disaster models, such as those of Barro and Backus et al., can explain the population level of volatility, but most of the volatility occurs during disaster periods. In fact, stock prices exhibit high level of volatility during disaster and non-disaster times.

Other recent work seeks to quantitatively explain implied volatilities in rare disaster models. Du (2011) extends the external habit model of Menzly, Santos, and Veronesi (2004) to allow for disasters that occur with constant probability. Our paper departs from his in that we do not require habit preferences and we allow the disaster probability to be stochastic.⁴ This paper also relates to that of Gabaix (2012), who incorporates time-varying disaster sensitivities into a model with additively separable utility using linearity-generating processes. In contrast, we focus on a stochastic probability of disasters and allow for recursive utility. Another point of departure from the earlier literature is our focus on the mechanism driving the level and slope of implied volatilities. Our framework directly nests earlier cases

⁴See also Nowotny (2011) who focuses on the implications of self-exciting processes for equity markets rather than on option prices.

considered by the literature (e.g. Backus, Chernov, and Martin (2011)), which allows us drill down to uncover the reason why we find different results. Moreover, besides the relation between implied volatilities and moneyness, we also examine the relation between implied volatilities and time to expiration. We show that our model can match this relation as well.

Our model suggests links between the probability of rare economic events, option prices, and equity risk premia. A growing line of work explores these links empirically. Bollerslev and Todorov (2011) show that a substantial fraction of the equity premium can be attributable to jump risk reflected in option prices, Gao and Song (2013) price crash risk in the cross-section using options, and Kelly, Pastor, and Veronesi (2014) demonstrate a link between options and political risk. The results in these papers provide empirical support for the theoretical mechanism that we highlight, namely that options reflect the risk of economy-wide rare events, and that this risk varies over time.

2 Model

This section describes our model. Section 2.1 details the primitives. Section 2.2 describes the pricing kernel and stock market values, with a focus on stock market volatility. Section 2.3 discusses a limiting case of our main model in which the probability of a disaster is constant. Section 2.4 describes the computation of option prices.

2.1 Model primitives

We assume an endowment economy with complete markets and an infinitely-lived representative agent. Aggregate consumption (the endowment) solves the following stochastic differential equation

$$\frac{dC_t}{C_t} = \mu dt + \sigma dB_t + (e^{Z_t} - 1) dN_t, \quad (1)$$

where B_t is a standard Brownian motion and N_t is a Poisson process with time-varying intensity λ_t . Following Wachter (2013), we assume the intensity follows the process

$$d\lambda_t = \kappa(\bar{\lambda} - \lambda_t) dt + \sigma_\lambda \sqrt{\lambda_t} dB_{\lambda,t}, \quad (2)$$

where $B_{\lambda,t}$ is also a standard Brownian motion, and B_t , $B_{\lambda,t}$ and N_t are assumed to be independent. For the range of parameter values we consider, λ_t is small and can therefore be interpreted to be (approximately) the probability that $dN_t = 1$. We thus will use the terminology probability and intensity interchangeably, while keeping in mind that the relation is an approximate one.⁵

In this model, a disaster is represented by an increment to the Poisson process, namely $dN_t = 1$. The size of the disaster is determined by the random variable $Z_t < 0$, which is the change in log consumption should a disaster occur. For convenience, we assume that Z_t has a time-invariant distribution which we will call ν . Outcomes of Z_t are independent of the other random variables. We will use the notation E_ν to denote expectations of functions of Z_t taken with respect to the ν -distribution.

The model above implies that the probability of a disaster varies over time according to a square-root process like that found in Cox, Ingersoll, and Ross (1985). It mean-reverts to $\bar{\lambda}$, and is therefore stationary. The presence of a square root implies that λ_t cannot be negative. Because λ_t is, strictly speaking, an intensity and not a probability, it is theoretically possible for λ_t to go above 1. While this presents no problems for the model, it is useful to note that this event is extraordinarily unlikely in our calibration.

We will assume a recursive generalization of power utility that allows for preferences over

⁵This model assumes that disasters occur instantaneously, while in the data, they are spread out over several years. Nakamura, Steinsson, Barro, and Ursúa (2013) and Tsai and Wachter (2015) show that implications for stock prices are similar to the instantaneous case provided that one assumes, as we do, that the agent has a preference for early resolution of uncertainty.

the timing of the resolution of uncertainty. Our formulation comes from Duffie and Epstein (1992), and we consider a special case in which the elasticity of intertemporal substitution is equal to one. That is, we define continuation utility V_t for the representative agent using the following recursion:

$$V_t = E_t \int_t^\infty f(C_s, V_s) ds, \quad (3)$$

where

$$f(C, V) = \beta(1 - \gamma)V \left(\log C - \frac{1}{1 - \gamma} \log((1 - \gamma)V) \right). \quad (4)$$

The parameter γ is relative risk aversion and the parameter β is the rate of time preference. This utility function is equivalent to the continuous-time limit (and the limit as the EIS approaches one) of the utility function defined by Epstein and Zin (1989) and Weil (1990).

2.2 The state-price density and the aggregate market

We first solve for the price of the claim to the aggregate dividend D_t . We make the standard assumption that $D_t = C_t^\phi$, for a parameter ϕ that is often referred to as leverage.⁶ We solve for asset prices using the marginal utility process for the representative agent, i.e. the state-price density.

Given our assumptions on the endowment and on utility, it follows that the state-price density is characterized by

$$\frac{d\pi_t}{\pi_t} = -(r_t + \lambda_t E_\nu [e^{-\gamma Z} - 1]) dt - \gamma \sigma dB_t + b\sigma_\lambda \sqrt{\lambda_t} dB_{\lambda,t} + (e^{-\gamma Z_t} - 1) dN_t. \quad (5)$$

⁶See e.g., Campbell (2003) and Abel (1999). This assumption implies that dividends respond more than consumption to disasters, an assumption that is plausible given the U.S. data (Longstaff and Piazzesi, 2004). The assumption of leverage is not crucial in that we would obtain similar results in a model with no leverage and a higher value of the EIS.

for

$$b = \frac{\kappa + \beta}{\sigma_\lambda^2} - \sqrt{\left(\frac{\kappa + \beta}{\sigma_\lambda^2}\right)^2 - 2\frac{E_\nu[e^{(1-\gamma)Z} - 1]}{\sigma_\lambda^2}}.$$

and where the riskfree rate r_t is equal to

$$r_t = \beta + \gamma\mu - \gamma\sigma^2 + \lambda_t E_\nu[e^{-\gamma Z}(e^Z - 1)]. \quad (6)$$

See Wachter (2013) and Appendix B.1 for a derivation.

Equation 5 summarizes how the shocks in the economy affect the marginal utility of the representative agent. The term $\gamma\sigma dB_t$ is what would appear in the CCAPM (Breedon, 1979). It shows how marginal utility is affected by normal-times consumption risk. For small values of γ , this term is negligible. The term $b\sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t}$ shows the effect of an increase in disaster risk λ_t . An increase in λ_t raises marginal utility as long as γ exceeds 1 (the EIS), as is apparent from the equation for b . That is, shocks to the probability of a disaster are priced. Finally, $(e^{-\gamma Z_t} - 1)dN_t$ represents the effect of the rare disaster itself. As is clear from the prior literature (Barro, 2006) this term is responsible for far more of the equity premium than is $\gamma\sigma dB_t$. Decreasing absolute risk aversion, as represented by isoelastic utility, implies that disaster realizations are far more costly to the investor than normal-times variation in consumption.

These assumptions imply that the stock price is a function of dividends D_t and the disaster probability λ_t . In equilibrium, the stock price must equal the sum (or, in the case of continuous time, the indefinite integral) of future dividends multiplied by the intertemporal marginal rate of substitution:

$$S(D_t, \lambda_t) = E_t \left[\int_t^\infty \frac{\pi_s}{\pi_t} D_s ds \right]. \quad (7)$$

Dividing by current dividends yields the price-dividend ratio:

$$G(\lambda_t) = E_t \left[\int_t^\infty \frac{\pi_s D_s}{\pi_t D_t} ds \right]. \quad (8)$$

Wachter (2013) shows that this expectation can be expressed in closed form with

$$G(\lambda_t) = \int_0^\infty \exp \{ a_\phi(\tau) + b_\phi(\tau) \lambda_t \}, \quad (9)$$

where

$$\begin{aligned} a_\phi(\tau) &= \left((\phi - 1) \left(\mu + \frac{1}{2} (\phi - \gamma) \sigma^2 \right) - \beta - \frac{\kappa \bar{\lambda}}{\sigma_\lambda^2} (\zeta_\phi + b \sigma_\lambda^2 - \kappa) \right) \tau \\ &\quad - \frac{2\kappa \bar{\lambda}}{\sigma_\lambda^2} \log \left(\frac{(\zeta_\phi + b \sigma_\lambda^2 - \kappa) (e^{-\zeta_\phi \tau} - 1) + 2\zeta_\phi}{2\zeta_\phi} \right) \\ b_\phi(\tau) &= \frac{2E_\nu [e^{(1-\gamma)Z_t} - e^{(\phi-\gamma)Z_t}] (1 - e^{-\zeta_\phi \tau})}{(\zeta_\phi + b \sigma_\lambda^2 - \kappa) (1 - e^{-\zeta_\phi \tau}) - 2\zeta_\phi}, \end{aligned}$$

and

$$\zeta_\phi = \sqrt{(b \sigma_\lambda^2 - \kappa)^2 + 2E_\nu [e^{(1-\gamma)Z_t} - e^{(\phi-\gamma)Z_t}] \sigma_\lambda^2}.$$

Because of the presence of rare disasters, this model can account for the well-known equity premium puzzle, the riskfree rate puzzle, and the level of stock market volatility.

Given our focus on option pricing, what is most relevant is the model's predictions for stock market volatility. Stock price variation can come from variation in dividends (cash flows), or from variation in the price-dividend ratio $G(\lambda_t)$. These sources are reflected in the law of motion for stock prices, which follows from $S_t = D_t G(\lambda_t)$ and Ito's Lemma:

$$\frac{dS_t}{S_{t-}} = \mu_{S,t} dt + \phi \sigma dB_t + \frac{G'(\lambda_t)}{G(\lambda_t)} \sigma_\lambda \sqrt{\lambda_t} dB_{\lambda,t} + (e^{\phi Z_t} - 1) dN_t, \quad (10)$$

where $\mu_{S,t}$ is the equilibrium-determined drift in the stock price. The term $\phi\sigma dB_t$ represents normal-times variation in dividends. Note that $\phi\sigma$ is dividend volatility. The term $(e^{\phi Z_t} - 1) dN_t$ represents the shock to dividends in case of a disaster. It is equal to $(e^{\phi Z_t} - 1)$ if a disaster occurs ($N_t = 1$) and zero otherwise. Leverage implies that dividends, and hence stock prices, respond more to disasters than does consumption.

Finally, the term $\frac{G'(\lambda_t)}{G(\lambda_t)}\sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t}$ represents how changes in the price-dividend ratio impact stock prices. Through the logic of Ito's Lemma, it is equal to the percent change in the price-dividend ratio coming from a change in λ_t , multiplied by the volatility of λ_t , which is equal to $\sigma_\lambda\sqrt{\lambda_t}$ (note that, because of the square root process, the volatility of stock prices rises in the level of disaster risk). As we will show below, this second term, the change in stock prices due to changes in the disaster probability, is the cause of almost all of the normal-times variation in stock prices. It is an order of magnitude greater than the variation due to dividends alone.

Why, economically, do changes in λ_t lead to changes in prices? Equation 10 implies that, when the risk of a disaster rises, the price-dividend ratio falls. The total effect on the price-dividend ratio can be decomposed into an effect on the riskfree rate, on the equity premium, and on expected cash flows. An increase in the risk of a disaster lowers the riskfree rate; this effect causes all stores of value to increase, including stock prices (holding all else equal). However, an increase in the disaster probability also increases the equity premium. The greater is the risk of a disaster, the more compensation the representative agent requires to hold equities. Finally, when the disaster risk is higher, expected future cash flows are lower. When leverage $\phi > 1$ (and assuming recursive utility) the second two effects dominate the first, and the price-dividend ratio is, realistically, decreasing in the probability of disaster.⁷

⁷There is direct evidence that stock prices fall with an increase in the disaster probability (Barro and Ursúa, 2009; Berkman, Jacobsen, and Lee, 2011; Manela and Moreira, 2015).

2.3 Constant disaster risk: A special case

A special case of the model above is that of constant disaster risk, namely $\lambda_t \equiv \bar{\lambda}$. This model is solved by Naik and Lee (1990) and is the continuous-time equivalent of the models of Barro (2006) and Backus, Chernov, and Martin (2011).⁸ Note that while these papers assume time-additive utility, this is observationally equivalent to recursive utility if consumption growth is iid. Indeed, the state-price density is characterized by

$$\frac{d\pi_t^{\text{CDR}}}{\pi_{t^-}^{\text{CDR}}} = -(r + \lambda E_\nu [e^{-\gamma Z_t} - 1]) dt - \gamma \sigma dB_t + (e^{-\gamma Z_t} - 1) dN_t, \quad (11)$$

where the riskfree rate r and the disaster probability λ are now constants.⁹

For this model with iid cash flows, the price-dividend ratio is a constant, and thus the equation for stock price fluctuations reduces to

$$\frac{dS_t^{\text{CDR}}}{S_{t^-}^{\text{CDR}}} = \mu_S dt + \phi \sigma dB_t + (e^{\phi Z_t} - 1) dN_t. \quad (12)$$

In this model, the only fluctuations in stock prices come from fluctuations in cash flows. The two terms in (12) denote normal-times and disaster-times respectively.

2.4 Option pricing

A European put option gives the holder the right to sell the underlying security at an expiration date T for an exercise price K . We take as our underlying the stock index described in the previous section. Because the payoff on the option is $(K - S_T)^+$, no-

⁸See Appendix D for the mapping from the discrete-time process in Backus, Chernov, and Martin (2011) to the continuous-time process that we use.

⁹Appendix B.2 derives the constant disaster risk case as a limit of SDR as $\sigma_\lambda \rightarrow 0$ and shows it is what one would find by deriving the CDR case from first principles. Appendix B.3 gives the isomorphism linking the recursive utility economy to the time-additive utility economy in the case of iid consumption.

arbitrage implies that

$$P(S_t, \lambda_t, T - t; K) = E_t \left[\frac{\pi_T}{\pi_t} (K - S_T)^+ \right].$$

Let $K^n = K/S_t$ denote the normalized strike price (“moneyness”) and $P_t^n = P_t/S_t$ the normalized put price. Then

$$P^n(\lambda_t, T - t; K^n) = E_t \left[\frac{\pi_T}{\pi_t} \left(K^n - \frac{S_T}{S_t} \right)^+ \right]. \quad (13)$$

While the price of a put option depends on S_t , the normalized price does not. The reason is that, at time t , λ_t contains all the information necessary to compute the distributions of $\frac{\pi_T}{\pi_t}$ and $\frac{S_T}{S_t}$. Because our ultimate interest is in implied volatilities, it suffices to compute normalized put prices.¹⁰

To calculate normalized put prices, we use the transform analysis of Duffie, Pan, and Singleton (2000), applied to a highly-accurate log-linear approximation of the price-dividend ratio. This analytical method avoids the need to simulate the expectation in (13) which is

¹⁰Given a stock price S , exercise price K , time to maturity $T - t$, interest rate r , and dividend yield y , the Black and Scholes (1973) put price is defined as

$$\text{BSP}(S, K, T - t, r, y, \sigma) = e^{-r(T-t)} KN(-d_2) - e^{-y(T-t)} SN(-d_1)$$

where $\Phi(\cdot)$ denotes the normal cumulative distribution function and

$$\begin{aligned} d_1 &= \frac{\log(S/K) + (r - y + \sigma^2/2)(T - t)}{\sigma\sqrt{T - t}} \\ d_2 &= d_1 - \sigma\sqrt{T - t}. \end{aligned}$$

Given normalized put prices, inversion of this Black-Scholes formula gives us implied volatilities. Specifically, the implied volatility $\sigma_t^{\text{imp}} = \sigma^{\text{imp}}(\lambda_t, T - t; K^n)$ solves

$$P_t^n(\lambda_t, T - t; K^n) = \text{BSP} \left(1, K^n, T - t, r_t^b, 1/G(\lambda_t), \sigma_t^{\text{imp}} \right)$$

where r_t^b is the model’s analogue of the Treasury Bill rate, which allows for a 40% probability of a default in case of a disaster (Barro, 2006).

inefficient due to rare events. See Appendix C.2 for details.

3 Option prices in the model and in the data

We now examine the implications of stochastic disaster risk for option prices. To ensure the model also matches equity data, we use the same parameters as Wachter (2013). Thus the parameters, given in Table 1, are chosen without appeal to option prices.¹¹ We use the disaster distribution of Barro and Ursúa (2008), who determine this distribution from international consumption data. Consumption declines of greater than 10% are said to be a disaster. The resulting distribution is multinomial, with a density function shown in Panel A of Figure 6.

We compute average implied volatilities in the SDR model by solving for implied volatilities as a function of λ_t (see Section 2.4), and then integrating over the stationary distribution.¹² For the CDR model, we use exactly the same parameters, except that we set $\sigma_\lambda = 0$. In this case, the model is equivalent to one with a constant probability of disaster equal to $\bar{\lambda}$. Below, we discuss the sensitivity of the results to our parameter choices.

Figure 2 shows three-month average implied volatilities in the data, in the SDR model, and in the CDR model. (see Appendix A for a data description). Implied volatilities are shown as a function of moneyness, namely exercise price divided by index level. Following Backus, Chernov, and Martin (2011), we consider a “moneyness” range that starts at 0.94, and, because of liquidity issues for in-the-money put options, ends at 1.02. As discussed in

¹¹We consider a parsimonious model with a single state variable. The resulting model can match the level and slope of the implied volatility curve, and can generate variation in the level over time. The slope of the curve will also vary, but to a lesser degree than in the data. To fully capture independent time-variation in the level and slope of the curve, a two-factor model is helpful (Seo and Wachter, 2016). However, the economic mechanism that allows the model to capture the most important stylized facts is the same in the more complicated model as compared with the simpler one.

¹²The stationary distribution is Gamma with shape parameter $2\kappa\bar{\lambda}/\sigma_\lambda^2$ and scale parameter $\sigma_\lambda^2/(2\kappa)$ (Cox, Ingersoll, and Ross, 1985). In the CDR model, implied volatilities are constant.

the introduction, implied volatilities in the data exhibit a skew: out-of-the-money (OTM) volatilities are 24% whereas at-the-money (ATM) volatilities are 21%. This skew rules out a conditional lognormal model. Of course, neither the CDR or the SDR model are lognormal: they put greater weight on extreme negative events. Moreover, the level of implied volatilities are higher than realized volatility, as shown in Figure 1.¹³

The dashed line in Figure 2 shows that the CDR model does generate a volatility skew. Intuitively, relative to a log-normal model, investors are willing to pay more for OTM options because they offer disaster insurance: they pay off when stock prices are low, and these states are more likely than they would be under a log-normal model. However, this volatility skew is strikingly different than in the data. While ATM volatility in the data is 21%, it is only 11% in the CDR model. Thus the level of the curve is much lower in the CDR model. However, the slope of the volatility curve is much higher. OTM volatility in the CDR model is 20%, far higher than the ATM volatility.

This striking failure of the CDR model to account for option prices in the data lead Backus, Chernov, and Martin (2011) to conclude that option prices are inconsistent with rare events as an explanation of the equity premium.¹⁴ Their interpretation is that investors do not view the international experience of the last 100 years as applicable to the postwar U.S. Instead, investors assume prices are subject to small frequent jumps, which can explain both option prices and the equity premium. Our next section offers an alternative and, we

¹³Given the rare-event nature of the models under discussion, one might wonder whether high implied volatility represents a Peso problem, such as that proposed to explain the equity premium (Brown, Goetzmann, and Ross, 1995; Jorion and Goetzmann, 1999). Perhaps the U.S. was unusually lucky, and volatility was lower than investors were expecting. To address this question, we simulate 10,000 alternative histories from the SDR model, with length equal to postwar data. The average realized volatility is 15%, about the same as in the data. Thus the finding of low realized volatility relative to implied volatility is typical.

¹⁴Backus, Chernov, and Martin (2011) trace out a very similar volatility curve but assume a different set of parameters. The difficulty with their exact choice of parameters, and why we do not use it, is that it implies non-convergence of prices (prices are a geometric series, and this series does not converge for certain parameter choices). In the CDR model, as in any iid model, prices are constant multiples of consumption, and so it is possible to miss this failure to converge. We will argue below that the slope of the curve implied by the CDR model is too steep under a wide range of parameter choices, all of which imply convergence.

argue, a more plausible interpretation of the data.

4 Discussion

The previous section shows that the SDR model can fit the options data, even with rare, large disasters. In this section, we reconcile this finding with the failure of the CDR model to fit options, even with the same disaster calibration. We also show that this finding holds across many reasonable calibrations, and we show further implications of both models, such as for the term structure of volatility.

Consider Table 1, which describes our benchmark parameters. All but the last two parameters apply to both the CDR and SDR models. These are: risk aversion γ , the EIS, the rate of time preference β , the average growth rate of consumption during normal times μ , the volatility of consumption growth during normal times σ , leverage ϕ , and the average probability of a rare disaster $\bar{\lambda}$. The volatility parameter of the λ_t process, σ_λ , is equal to zero by definition in the CDR model, and the mean reversion in the disaster probability, κ , is undefined.

In the CDR model, the EIS and β are not separately identified (see Appendix B.3), though they are in the SDR model. Moreover, changes in β and μ have only second-order effects on option pricing, through the level of the dividend yield. In Sections 4.1-4.4, we focus on the CDR model, and thus discuss the effects of γ , σ , ϕ , and $\bar{\lambda}$.

4.1 The role of leverage

One parameter that has been implicated in the failure of the CDR model to match options data is ϕ , interpreted as financial leverage.¹⁵ Leverage is an interesting candidate because

¹⁵The original calibration of Backus, Chernov, and Martin (2011) has a much higher leverage. Gabaix (2012) conjectures that this may be why they report a steep slope.

it increases the risk of equity during disaster times; the greater is ϕ the more responsive dividends are to a consumption decline. Intuitively, a lower value of leverage might imply a shallower slope of the implied volatility curve, bringing it more in line with the data.

Figure 3 compares the CDR model with one with lower leverage. That is, rather than leverage of 2.6, we consider leverage of 1.6. In fact, lowering leverage mainly reduces the level of the curve, with very little change in the slope. Because the level of implied volatilities in the CDR model is already too low given the data, reducing leverage brings the model further away from the data.

There are two reasons for this perhaps counterintuitive result. The first is that reducing leverage also reduces normal-times volatility. As we will see, the level of the curve is determined in large part by normal-times volatility, namely the volatility of the stock price process outside of the times when disaster is taken place. While much of the previous literature on consumption-based asset pricing has focused on generating correct volatilities in calibrated models – essentially matching population volatilities to measured volatilities in the data – less attention has been paid to the conditional distribution of volatility, namely “when” the volatility occurs. In part this is because conditional distributions are difficult to measure given stock return data alone. By providing a relation between implied volatilities and moneyness, option prices give information on this conditional distribution.

The CDR model implies that price changes have a volatility of 17%, similar to that of the data. However, this is the population value of price changes, most of the time what will be observed is the much lower volatility of normal-times shocks to dividends (recall that in the CDR model, the volatility of prices and of dividends are equal). This normal-times volatility is equal to $\phi\sigma = 2.6 \times 2 = 5.12\%$. The fact that normal-times volatility is much lower than the population volatility is reflected in a value of ATM volatility that is far below 17%.¹⁶ Investors are not willing to pay much to insure against small shocks under this model.

¹⁶Yan (2011) shows analytically that, as the time to expiration approaches zero, the implied volatility for

Lowering leverage compounds the problem because it further reduces normal-times volatility.

There is a second reason why reducing leverage mainly reduces the level of the curve. That is, even the price of ATM options depend on the disaster distribution. This makes sense, as these instruments also act as a hedge against disaster. Because both of these effects operate, the effect on ATM options of lowering leverage is very slightly larger than the effect of on OTM options, leading to a slightly higher slope. Clearly changing leverage cannot bring the CDR model in line with the data.

4.2 The role of risk aversion

A second parameter that determines implied volatilities is risk aversion. Risk aversion stands between the physical distribution of the consumption process and the risk-neutral distribution, as reflected in implied volatilities. Put options are hedges against the occurrence of a disaster. The greater is risk aversion, the more investors are willing to pay for this hedge. This should raise put prices, and also implied volatilities.

Indeed, Figure 3 shows that increasing relative risk aversion from 3 to 4 in the CDR model has a dramatic effect on implied volatilities. The ATM implied volatility is 17% rather than 11% (the data value is 21%). However, implied volatilities increase on OTM options by a still greater amount. The model value is 27%, rather than 20%, above the data value of 24%. Thus, while increasing risk aversion can help solve the problem with ATM volatilities, it does so at the cost of making OTM volatilities too high. Like leverage, risk aversion also mainly brings about a change in the level of the curve, not the slope, though there is a small effect on the slope as well. Thus changing risk aversion alone cannot bring the CDR model in line with the data.

ATM options is equal to the normal-times volatility in the stock price. However, with a time-to-expiration of three months, rare events do affect the ATM volatility. Otherwise it would be 5% rather than 11%.

4.3 The role of the disaster probability

We now turn to the probability of disaster. We consider the effect of increasing the probability by two percentage points to 5.55%. This probability of disaster is not empirically reasonable, but that does not matter for the thought exercise.

As one might expect, raising the disaster probability increases the implied volatilities for OTM options. It increases the volatility for ATM options too, but by slightly less. Thus, like leverage and risk aversion, the largest effect of raising the disaster probability is on the level, not the slope of the curve.

If we were to raise the disaster probability much more (so that the disaster interpretation of the Poisson shocks were called into question), while at the same time reducing the size of disasters, we would begin to see an increase in the level and a flattening of the curve because the resulting distribution would become more normal. This is in fact the solution that Backus, Chernov, and Martin (2011) propose for matching implied volatilities. Our results show that this solution is not necessary for matching implied volatilities; stochastic disaster risk is an alternative. Moreover, frequent, small disaster can be ruled out by consumption data itself, as discussed in Section 4.9.

4.4 The role of normal-times consumption volatility

The previous discussion demonstrated that changes in leverage or risk aversion, or small changes in the disaster probability, fail to bring the CDR model close to the data. The intuition we develop suggests that, to bring the CDR model more in line with the data, it is necessary to have a mechanism that specifically increases volatility during non-disaster periods. This interpretation of the failure of the CDR model is quite different from that in Backus, Chernov, and Martin (2011); the problem is not that the disasters are too large, but rather that the normal-times volatility is too small.

Panel A of Figure 4 considers the benchmark CDR model with a normal-times consumption volatility parameter $\sigma = 2\%$, and an alternative calibration of $\sigma = 4\%$. Consistent with the above intuition, ATM volatilities rise substantially under the alternative calibration, from 11% to 15%. OTM volatilities rise as well, but by much less, from 19.6% to 20.3%. It seems that raising the level of normal-times consumption volatility has promise to bring the CDR model in line with the data.

For this method to work, however, one would have to reconcile this higher σ with the very low observed volatility in postwar data. Following Barro (2006), the benchmark calibration of the CDR and SDR models assume $\sigma = 2\%$, the average postwar volatility of output in G7 countries. We ask: supposing that the CDR model with $\sigma = 4\%$ were correct, what is the probability of observing a volatility as low as 2%?

To answer this question, we simulate 10,000 samples of length designed to match the postwar data using the CDR model with $\sigma = 4\%$. The results are shown in Panel B. The distribution has a long right tail, consistent with the presence of rare disasters. The modal value is 4%, consistent with this being normal-times volatility. The left tail, in contrast, is extremely thin. The minimum observation across the 10,000 samples is still well above 2%, implying that observing a consumption volatility this low is astronomically unlikely. Nonetheless, Panel A shows that a σ as high as even 4% is insufficient to explain the level and slope of the implied volatility curve.¹⁷

While raising normal-times consumption volatility, by itself, cannot resolve the differences between the CDR model and the data, it reduces the difference considerably. It does so by increasing OTM and ATM volatilities, but it increases ATM volatilities by far more. However, raising implied volatility by increasing the volatility of cash flows is inconsistent with the data, for the same reason that stock market realized volatility cannot be explained

¹⁷This calibration implies a volatility of log dividend growth of 10.4% during normal times, which is also far above the 6.5% observed in postwar data.

by cash flow volatility alone (Shiller, 1981). The volatility in post-war cash flow data is simply too small.

Given the importance of normal-times volatility, a natural question is whether the success of the SDR model relies on a value for consumption volatility of 2%. While 2% may be an average taken across G7 countries, in the U.S., the average volatility of consumption in the U.S. is 1.5%. Perhaps with this lower volatility, the SDR model would fail to match the data. However, reducing σ all the way to 1% (which would imply a normal-times consumption volatility that is lower than in the post-war data) has a barely discernable effect on the implied volatility curve of the SDR model, as shown in Figure 5. There are two reasons for this. First, lowering normal-times volatility has a smaller effect than raising it, even in the CDR model. ATM options insure not only against normal-times variation, but also against rare disasters. As normal-times volatility falls, it has less of an impact on the price of the option and hence on implied volatilities. Second, in the SDR model, volatility in cash flows is a relatively small part of volatility in prices. Unlike the CDR model where cash flow volatility is the unique source of volatility in prices, prices change in the SDR model mostly because of changes in the disaster probability.¹⁸ While raising the volatility of consumption makes the CDR model look somewhat more like the SDR model, it is not the case that lowering the volatility of consumption makes the SDR model more like the CDR model.

4.5 The role of the disaster distribution

We now ask about the robustness of our calibration of the SDR model. One concern, voiced by Chen, Dou, and Kogan (2013), is that disaster models are sensitive to inherently unobservable features of the data (“dark matter”), for example, the disaster distribution.

¹⁸To be precise, total normal-times return volatility in the SDR model equals the square root of the variance due to λ_t , plus the variance in dividends. Dividend variance is small, and it is added to something much larger to determine total variance. Thus the effect of dividend volatility on return volatility is very small, and changes in dividend volatility also have relatively little effect.

We check the robustness of our results by considering an alternative distribution of disasters, proposed by Backus, Chernov, and Martin (2011). Like our benchmark distribution, this distribution can also explain the equity premium under the CDR model. Unlike a multinomial model, this distribution is lognormal, with mean parameter 30% and standard deviation 15%. We show the probability density function of both distributions in Panel A of Figure 6. They are very different. The multinomial distribution has a large mass around relatively small disasters (10-20% consumption declines), and then very thick tail (there is some mass in the 50-70% region). The lognormal distribution peaks at 30%, and then has a much thinner tail. Nonetheless, Panel B shows that they imply very similar option prices, both for the CDR model and the SDR model.¹⁹

4.6 The role of preferences for early resolution of uncertainty

We now turn to the role played by variation in the disaster probability. One way this variation could influence option prices is through the representative agent's preference for early resolution of uncertainty.

The SDR model assumes utility of the Epstein and Zin (1989) form, a generalization of time-additive power utility that allows for a preference over the resolution of uncertainty. Our assumptions of risk aversion γ equal to 3 and an EIS equal to 1 imply that the agent prefers an early resolution of uncertainty, in effect creating aversion to time-variation in λ_t .

A preference for an early resolution of uncertainty could make implied volatilities higher than otherwise, contributing to the success of the SDR model. As shown in Section 2.2, the marginal utility of the representative agent rises when the disaster probability rises; namely shocks to the disaster probability are priced. Prices of put options are an increasing function

¹⁹In a similar vein, we have also computed option prices using the power law distribution proposed by Barro and Jin (2011) with a tail parameter of 6.5. The results are indistinguishable from those of our benchmark calibration.

of the probability of a disaster. Therefore put options will have lower expected returns, and higher average prices, than if this effect were absent. It follows that implied volatilities will also be higher, since the option price is an increasing function of implied volatility. Even though the CDR model also assumes recursive utility, this mechanism is absent because consumption growth is iid, and the pricing kernel is observationally equivalent to one with time-additive utility (Appendix B.3).

To evaluate the quantitative relevance of this mechanism, we compute option prices assuming that the representative agent has time-additive (power) utility, namely no preference for early resolution of uncertainty. Because we want to isolate the specific effect of recursive utility on option prices, we continue to assume that stock prices follow the process (10), with parameters identical to what was assumed in our benchmark case.

The power utility investor has utility

$$V_{p,t} = \int_t^\infty e^{-\beta s} \frac{C_s^{1-\gamma}}{1-\gamma} ds.$$

The state-price density is simply

$$\pi_{p,t} = \beta^t C_t^{-\gamma}, \tag{14}$$

where we use the subscript p to denote power utility. From Ito's Lemma it follows that

$$\frac{d\pi_{p,t}}{\pi_{p,t^-}} = -(r_{p,t} + \lambda_t E_\nu [e^{-\gamma Z_t} - 1]) dt - \gamma \sigma dB_t + (e^{-\gamma Z_t} - 1) dN_t. \tag{15}$$

Equation 15 takes the same form as π_t in the recursive utility case, (5), except that b is set to zero.²⁰ This is intuitive, because b determines the pricing of shocks to the disaster

²⁰Under power utility, the riskfree rate is different, namely

$$r_{p,t} = \beta + \gamma \mu - \frac{1}{2} \gamma (\gamma + 1) \sigma^2 - \lambda_t E_\nu [e^{-\gamma Z_t} - 1]. \tag{16}$$

probability. For a given λ_t , it is also the same as the CDR pricing kernel (11) because, by definition, there is no pricing of λ_t -risk when λ_t is constant. To summarize, we compute normalized put prices, (13), with pricing kernel given by (14), and stock prices given by (10).

Figure 7 shows the implied volatilities (line with triangles) for the power utility case. For comparison, the figure also shows the benchmark SDR case, as well as the CDR case. As the theoretical discussion suggests, the power utility pricing kernel implies lower implied volatilities than the benchmark case. The difference, however, is not large, especially when compared with the difference between the SDR model and the CDR model. For ATM options, the power utility pricing kernel gives an implied volatility 18%, compared with 20% in the benchmark case. The CDR implied volatility is 11%. The preference for early resolution of uncertainty explains less than one quarter of the total difference between the CDR and SDR models. Thus, while some of the model's success is due to the pricing of time-varying disaster risk, it is not the main reason why the model is able to match the data.

The discussion so far has fixed the process for stock prices to that implied by recursive utility, with the purpose of focusing the discussion of the effect of recursive utility on options, over and above its effect on stocks. However, the assumption of recursive utility plays a substantial role in generating a process for stock returns that resembles the data. Coming back to the sources of stock-price volatility discussed in Section 2.2, an increase in the probability of disaster affects the riskfree rate as well as the equity premium and future cash flows. Depending on the parameters of the model, it is possible that the first effect will outweigh the second. Besides delivering the counterfactual implication that prices increase when the probability of a disaster rises, this would imply a too-low equity premium and

To compute option prices in the power utility case, we assume (14), which implies (16). Alternatively, we could have kept the riskfree rate the same as in the recursive utility case; namely, replacing (16) with (6) in the power utility pricing kernel. The two alternatives yield virtually indistinguishable implied volatilities.

highly volatile riskfree rates. In our model with recursive utility, the weak condition that $\phi > 1$ rules out this behavior. A model with power utility would require a much higher leverage, $\phi > \gamma$, to generate stock prices that are decreasing in the disaster probability; even then it is possible that riskfree rates would be too volatile.²¹ Clearly, recursive preferences are needed to obtain reasonable behavior of stock prices. Given reasonable stock prices, however, recursive utility plays a relatively minor role in option pricing.

4.7 The role of stock price volatility in normal times

The previous sections show that option prices in the SDR model are not determined by normal-times consumption volatility, nor by the precise form of the disaster distribution. Moreover, the pricing of risk of time-varying λ_t (which arises from a preference for early resolution of uncertainty) does not play a major role. What allows the SDR model to match option prices is a high normal-times volatility of stock prices, combined with rare disasters. When the disaster probability rises, stock prices fall, as shown in (10), and discussed in Section 2.2. Thus stock prices can fluctuate substantially even in the absence of large fluctuations in fundamentals.

This normal-times fluctuation in stock prices is crucial for matching the level and slope of the implied volatility curve. Sufficient normal-times volatility in stock prices generates high implied volatilities for ATM options. In contrast, while the CDR model can match the unconditional volatility in the market, most of this volatility occurs during periods with disasters. Thus ATM implied volatilities are low, and, because of the importance of disaster risk, the curve is very steep.

The success in matching implied volatilities comes about simply by making a constant

²¹Furthermore, for similar reasons, such a model would not be robust to the realistic change of allowing disasters to unfold over a finite length of time rather than instantaneously (Nakamura, Steinsson, Barro, and Ursúa, 2013; Tsai and Wachter, 2015).

parameter stochastic. Even though average implied volatilities are an *unconditional* moment of the data, a change in the *conditional* distribution in the model has a large effect. The results are in contrast to the impact of a similar exercise in a reduced-form model. In writing about whether introducing stochasticity would impact the curve, Backus, Chernov, and Martin (2011) write:

The question is whether the kinds of time dependence we see in asset prices are quantitatively important in assessing the role of extreme events. It is hard to make a definitive statement without knowing the precise form of time dependence, but there is good reason to think its impact could be small. The leading example in this context is stochastic volatility, a central feature of the option-pricing model estimated by Broadie et al. (2007). However, average implied volatility smiles from this model are very close to those from an iid model in which the variance is set equal to its mean. Furthermore, stochastic volatility has little impact on the probabilities of tail events, which is our interest here.

The above quote is correct that allowing the probability of disaster to vary over time has little impact on the unconditional probability of rare events. The average probability of disaster is the same in both models. Should a disaster occur, price changes are driven almost exclusively by changes in economic fundamentals, which have the same distribution in both models. By the logic of this paragraph, then, there should be little effect on option prices.

The intuition in this paragraph comes from the long tradition of reduced-form modeling in option pricing. Reduced-form studies such as Bates (2000), Broadie, Chernov, and Johannes (2007), Eraker (2004), Pan (2002), and Santa-Clara and Yan (2010)), emphasize the need for non-conditional-normality to explain implied volatilities. Like Backus, Chernov, and Martin (2011), these papers find a role for small, frequent jumps in stock prices. In the reduced-form literature, the process for stock prices S_t and for the pricing kernel π_t are exogenously

specified. That is, these models determine the properties of S_t and π_t directly to match the data. What they do not do is connect these properties back to the fundamentals of the economy.

In an equilibrium model, by contrast, allowing for time-variation in λ_t has a first-order effect. The reason is that the conditional properties of the disaster probability distribution affect the unconditional moments of stock returns. Unconditional stock market volatility arises endogenously from conditional moments of fundamentals, as can be seen from the equation for the stock price dynamics, (10). As this equation shows, time-variation in λ_t is the reason for normal-times variation in the stock price, a result that arises endogenously from investors preferences and expectations about cash flows (in a reduced-form model, normal-times variation in stock prices would be characterized by a separate shock from an unmodeled source). While in the reduced-form literature, the difference between iid and dynamic models principally affects conditional moments, in the equilibrium literature, the difference can affect the level of volatility itself.

4.8 The term structure of volatility

The previous discussion focused on three-month options. What about the term structure of volatility? Figure 8 shows ATM volatilities as a function of time to expiration. The figure shows 1, 3, and 6-month volatilities. We see from this figure that the SDR model can fit implied volatilities, not just at the 3-month horizon, but at other horizons as well. That is, the slope of the volatility term structure in the data and the model are quite close.

Figure 8 shows that both the SDR model and the data exhibit a slightly upward sloping term structure for ATM volatilities. While one might expect the CDR model, given that it is iid, to exhibit a flat term structure, it too exhibits an upward slope, though the level is quite different. In fact, the upward slope is slightly more pronounced: the difference between the

6-month and the 1-month volatility for the CDR model is 5.4%, while for the SDR model it is 3.9%. Why is this?

In an lognormal model, the implied volatilities would be independent of the horizon. However, in an iid model with rare disasters, the central limit theorem acts to bring the distribution of log price changes over longer horizons closer to normal (the log of the six month return is the sum of six one-month returns). That is, the six-month distribution for the CDR model looks (a bit) more like the distribution for a lognormal model with a population volatility of 17%.²² We can see this in the skewness of log price changes; this rises from -30 at the 1-month horizon to -12 at the six-month horizon. As the distribution of price changes becomes more normal, ATM prices increase.

The same forces are at work for the SDR model. The skewness rises from -13 at the one-month horizon to the -5 at the six-month horizon. Returns are substantially less skewed in the SDR model because of the higher level of normal-times volatility, and this skewness does not fall as quickly because of persistence in the disaster probability.²³

To summarize, the model matches the slight upward slope in ATM implied volatilities. It does so because of the central limit theorem: long-horizon returns look more normal than short-horizon returns. The persistence of disaster risk implies that the central limit theorem works more slowly than in an iid model, which helps bring the model in line with the data.

4.9 Alternative mechanisms

While it is not the purpose of this article to rule out all other potential explanations, this section briefly discuss alternative models that have been used to address option prices.

²²Neuberger (2012) also points out that the decay of skewness in an iid model is too fast compared with that implied by options data He argues that this is a reason to reject iid models..

²³The skewness in the SDR and CDR models arise from the presence of disasters. In our sample, we do not see this skewness in the data, but we do not see (large) disasters either. Skewness in the data is therefore well within standard error bands implied by the model.

Backus, Chernov, and Martin (2011) propose one such alternative mechanism, namely, that the consumption growth distribution is characterized by declines that are smaller and more frequent. This distribution is consistent with average implied volatilities as well as with the equity premium, and the mean and volatility of consumption growth observed in the U.S. in the 1889-2009 period (provided a coefficient of relative risk aversion equal to 8.7). However, this consumption distribution can be ruled out based on the history of consumption itself. Because it assumes that negative consumption jumps are frequent (as they must be to explain the equity premium), many would have occurred in the postwar period in the U.S. The unconditional volatility of consumption growth in the U.S. during this period was less than 2%. Under the option-implied consumption growth distribution, there is less than a 1 in one million chance of observing a 70-year period with volatility this low.

What about alternative modifications to the consumption distribution? In light of the discussion in Section 4.7, one such modification would be to allow volatility in consumption to be stochastic. This could contribute to stock price volatility. However, consumption volatility does not appear to vary enough to explain equity volatility, nor does the resulting economy deviate sufficiently from unconditional normality to explain the level of the implied volatility curve.²⁴ The existing literature has explored rich models of consumption dynamics that include time-varying volatility (see Benzoni, Collin-Dufresne, and Goldstein (2011), Buraschi and Jiltsov (2006) and Drechsler (2013)). These papers find that, while the mechanism can capture deviations from the Black-Scholes benchmark, it is insufficient to match the extent of the deviation as measured by the skew in the data.²⁵

As discussed in the introduction, simply explaining normal-times volatility is insufficient: the models of Campbell and Cochrane (1999) or Bansal and Yaron (2004) do not generate a

²⁴See Lettau, Ludvigson, and Wachter (2008); the persistent shifts in consumption volatility present in the data can explain changes in the level of stock prices, but not month-to-month variation in prices.

²⁵For example, the model of Benzoni, Collin-Dufresne, and Goldstein (2011) has a risk aversion of 10, but generates ATM implied volatilities of 16% compared with a data value of 21%.

volatility skew, even though they match the volatility of stock returns. Du (2011), Benzoni, Collin-Dufresne, and Goldstein (2011) and Drechsler and Yaron (2011) show that adding non-normalities to these models helps in fitting options data. However, in these cases, rare events are required as an additional mechanism in models already designed to match the equity premium and volatility.

Other models that can quantitatively explain implied volatilities do so by making non-standard assumptions on utilities or beliefs. For example, Drechsler (2013) assumes ambiguity aversion and Shaliastovich (2015) assumes jumps in confidence. These papers build on earlier work (Bates, 2008; Liu, Pan, and Wang, 2005) that shows that crash aversion or ambiguity aversion is necessary to reconcile option prices and equity prices in the context of an iid model. One way to characterize this literature is that models that can explain the equity premium (or, in the case of dynamic models, the equity premium and volatility) can have difficulty explaining options without the addition of non-standard preferences or beliefs. The present model is an exception.

5 Conclusion

The well-known implied volatility skew on equity index options has long constituted important evidence for excess kurtosis in stock returns. Separately, a literature has developed proposing excess kurtosis, modeled as rare consumption disasters, as an explanation for the equity premium puzzle. Relating the volatility skew and the equity premium is an important challenge because the same model that explains stock market returns should also explain option prices. However, standard models indicate that the small jumps in stock prices needed to match options data are qualitatively different from the rare disasters required to reconcile the equity premium with low observed consumption volatility. This seeming contradiction raises questions about the source of the equity premium, the volatility skew, and nature of

expectations of rare events in asset markets.

We have proposed an alternative and more general approach to modeling the risk of downward jumps that can reconcile the volatility skew and the equity premium. Rather than assuming that the probability of a large negative event is constant, we allow it to vary over time. Under our model, the existence of consumption disasters leads to both a volatility skew and an equity premium. Moreover, the variation in the probability of these events raises the level of at-the-money volatilities and generates the excess volatility observed in stock prices. Thus the model can simultaneously match the equity premium, equity volatility, and implied volatilities on index options. Option prices, far from ruling out rare consumption disasters, provide additional evidence for their importance in asset pricing.

Appendix

A Data construction

Our sample consists of daily data on option prices, volume and open interest for European put options on the S&P 500 index from OptionMetrics. Data are from 1996 to 2012.²⁶ Options expire on the Saturday that follows the third Friday of the month. We extract monthly observations using data from the Wednesday of every option expiration week. We apply standard filters to ensure that the contracts on which we base our analyses trade sufficiently often for prices to be meaningful. That is, we exclude observations with bid price smaller than 1/8 and those with zero volume and open interest smaller than one hundred contracts (Shaliastovich, 2015).

OptionMetrics constructs implied volatilities using the formula of Black and Scholes (1973) (generalized for an underlying that pays dividends), with LIBOR as the short-term interest rate. The dividend-yield is extracted from the put-call parity relation. We wish to construct a data set of implied volatilities with maturities of 1, 3 and 6 months across a range of strike prices. Of course, there will not be liquid options with maturity precisely equal to, say, 3 months, at each date. For this reason, we use polynomial interpolation across strike prices and times to expiration.²⁷ Specifically, at each date in the sample, we regress implied volatilities on a polynomial in strike price K and maturity T :

$$\sigma(K, T) = \theta_0 + \theta_1 K + \theta_2 K^2 + \theta_3 T + \theta_4 T^2 + \theta_5 KT + \theta_6 KT^2 + \epsilon_{K,T}$$

²⁶OptionMetrics data therefore does not cover the period prior to the 1987 crash, during which the implied volatility skew was reportedly smaller. Explaining the reason for the change around 1987 is beyond the scope of this paper. One possibility is that options were mechanically and incorrectly priced prior to 1987 using the Black and Scholes (1973) model (Constantinides, Jackwerth, and Perrakis, 2009).

²⁷See Dumas, Fleming, and Whaley (1998), Christoffersen and Jacobs (2004) and Christoffersen, Heston, and Jacobs (2009).

We run this regression on options with maturities ranging from 30 to 247 days, and with moneyness below 1.1. The implied volatility surface is generated by the fitted values of this regression.

B The state-price density

B.1 The state-price density in the SDR model

Duffie and Skiadas (1994) show that the state-price density π_t equals

$$\pi_t = \exp \left\{ \int_0^t \frac{\partial}{\partial V} f(C_s, V_s) ds \right\} \frac{\partial}{\partial C} f(C_t, V_t). \quad (\text{B.1})$$

Equation (B.1) shows the state-price density can be expressed in terms of a locally deterministic term and a term that is locally stochastic. We require both to be expressed in terms of C_t and λ_t . We derive the result for the stochastic term first.

It follows from (4) that

$$\frac{\partial}{\partial C} f(C_t, V_t) = \beta (1 - \gamma) \frac{V_t}{C_t}. \quad (\text{B.2})$$

Wachter (2013) shows that continuation utility V_t can be expressed in terms of C_t as follows:

$$V_t = J(\beta^{-1}C_t, \lambda_t), \quad (\text{B.3})$$

where $\beta^{-1}C_t = W_t$ is the wealth of the representative agent (because the EIS is equal to 1, β is the consumption-wealth ratio). The function J solves

$$J(W_t, \lambda_t) = \frac{W_t^{1-\gamma}}{1-\gamma} e^{a+b\lambda_t}, \quad (\text{B.4})$$

and

$$a = \frac{1-\gamma}{\beta} \left(\mu - \frac{1}{2}\gamma\sigma^2 \right) + (1-\gamma) \log \beta + b \frac{\kappa\bar{\lambda}}{\beta} \quad (\text{B.5})$$

$$b = \frac{\kappa + \beta}{\sigma_\lambda^2} - \sqrt{\left(\frac{\kappa + \beta}{\sigma_\lambda^2} \right)^2 - 2 \frac{E_\nu [e^{(1-\gamma)Z_t} - 1]}{\sigma_\lambda^2}}. \quad (\text{B.6})$$

For future reference, we note that b is a solution to the quadratic equation

$$\frac{1}{2}\sigma_\lambda^2 b^2 - (\kappa + \beta)b + E_\nu [e^{(1-\gamma)Z_t} - 1] = 0. \quad (\text{B.7})$$

Substituting (B.3) and (B.4) into (B.2) implies that

$$\frac{\partial}{\partial C} f(C_t, V_t) = \beta^\gamma C_t^{-\gamma} e^{a+b\lambda_t} \quad (\text{B.8})$$

It also follows from (4) that

$$\frac{\partial}{\partial V} f(C_t, V_t) = \beta(1-\gamma) \left(\log C_t - \frac{1}{1-\gamma} \log((1-\gamma)V_t) \right) + \beta$$

Substituting in for V_t from (B.3) and (B.4) implies

$$\frac{\partial}{\partial V} f(C_t, V_t) = \beta(1-\gamma) \log \beta - \beta(a + b\lambda_t) - \beta \quad (\text{B.9})$$

Finally, we collect constant terms. Define

$$\eta = \beta(1-\gamma) \log \beta - \beta a - \beta \quad (\text{B.10})$$

so that

$$\frac{\partial}{\partial V} f(C_t, V_t) = \eta - \beta b \lambda_t$$

Therefore, from (B.1) it follows that the state-price density can be written as

$$\pi_t = \exp\left(\eta t - \beta b \int_0^t \lambda_s ds\right) \beta^\gamma C_t^{-\gamma} e^{a+b\lambda_t} \quad (\text{B.11})$$

Equation 5 follows from Ito's Lemma; see Wachter (2013) for details.

B.2 The iid limit of the SDR model.

In this section we compute the limit of the state price density π_t . Note that b in equation (B.6) can be rewritten as

$$b = \frac{1}{\sigma_\lambda^2} \left(\kappa + \beta - \sqrt{(\kappa + \beta)^2 - 2E_\nu [e^{(1-\gamma)Z_t} - 1] \sigma_\lambda^2} \right).$$

Applying L'Hopital's rule implies that

$$\begin{aligned} \lim_{\sigma_\lambda \rightarrow 0} b &= \lim_{\sigma_\lambda \rightarrow 0} \frac{1}{2} \left((\kappa + \beta)^2 - 2E_\nu [e^{(1-\gamma)Z_t} - 1] \sigma_\lambda^2 \right)^{-\frac{1}{2}} 2E_\nu [e^{(1-\gamma)Z_t} - 1] \\ &= \frac{E_\nu [e^{(1-\gamma)Z_t} - 1]}{\kappa + \beta} \end{aligned}$$

It follows from the equation for a , (B.5), that

$$\begin{aligned} \lim_{\sigma_\lambda \rightarrow 0} (a + b\lambda_t) &= \lim_{\sigma_\lambda \rightarrow 0} (a + b\bar{\lambda}) \\ &= \frac{1-\gamma}{\beta} \left(\mu - \frac{1}{2}\gamma\sigma^2 \right) + (1-\gamma) \log \beta + (\kappa + \beta) \frac{\bar{\lambda}}{\beta} \lim_{\sigma_\lambda \rightarrow 0} b \\ &= \frac{1-\gamma}{\beta} \left(\mu - \frac{1}{2}\gamma\sigma^2 \right) + (1-\gamma) \log \beta + \frac{E_\nu [e^{(1-\gamma)Z_t} - 1] \bar{\lambda}}{\beta} \end{aligned}$$

where we assume that $\lambda_0 = \bar{\lambda}$ and therefore that $\lambda_t = \bar{\lambda}$ for all t .

We consider the limit of π_t/π_0 as σ_λ approaches zero. Note that for computing asset

prices, we only care about ratios of the state price density at different points of time, and so it suffices to compute this quantity. It follows from (B.1), (B.8) and (B.9) that

$$\begin{aligned} \lim_{\sigma_\lambda \rightarrow 0} \frac{\pi_t}{\pi_0} &= \exp \left\{ \left(\beta(1-\gamma) \log \beta - \beta - \beta \lim_{\sigma_\lambda \rightarrow 0} (a + b\bar{\lambda}) \right) t \right\} \left(\frac{C_t}{C_0} \right)^{-\gamma} \\ &= \exp \left\{ \left(-\beta - (1-\gamma) \left(\mu - \frac{1}{2} \gamma \sigma^2 \right) - E_\nu [e^{(1-\gamma)Z_t} - 1] \bar{\lambda} \right) t \right\} \left(\frac{C_t}{C_0} \right)^{-\gamma}, \end{aligned}$$

which is equivalent to the result one obtains by calculating the state price density in the iid case when the EIS is equal to 1. For this result to hold, it is important that we choose the lower of the two roots in (B.7), as pointed out by Tauchen (2005).

B.3 An isomorphism with power preferences under the iid assumption

In this section we show that, in an iid model, ratios of the state price density at different times implied by power utility are the same as those implied by recursive utility assuming the discount rate is adjusted appropriately. Thus the power utility model and the recursive utility model are isomorphic when the endowment process is iid.

Let $\pi_{p,t}$ be the state price density assuming power utility with discount rate β_p and relative risk aversion γ . Then

$$\frac{\pi_{p,t}}{\pi_{p,0}} = e^{-\beta_p t} \left(\frac{C_t}{C_0} \right)^{-\gamma}.$$

For convenience, let π_t be the state price density for recursive utility (with EIS equal to one).

As shown in Appendix B.2,

$$\frac{\pi_t}{\pi_0} = e^{((1-\gamma)(-\mu + \frac{1}{2} \gamma \sigma^2) - \bar{\lambda} E_\nu [e^{(1-\gamma)Z_t} - 1] - \beta) t} \left(\frac{C_t}{C_0} \right)^{-\gamma}. \quad (\text{B.12})$$

It follows that, for β given by

$$\beta = \beta_p + (1 - \gamma) \left(-\mu + \frac{1}{2}\gamma\sigma^2 \right) - \bar{\lambda} E_\nu [e^{(1-\gamma)Z_t} - 1],$$

ratios of the state price densities are the same.

C Details of the calculation of option prices

C.1 Approximating the price-dividend ratio

The formula for the price-dividend ratio in the SDR model is derived by Wachter (2013) and is given by

$$G(\lambda_t) = \int_0^\infty \exp \{a_\phi(\tau) + b_\phi(\tau)\lambda_t\} d\tau,$$

where $a_\phi(\tau)$ and $b_\phi(\tau)$ have closed-form expressions given in that paper. The algorithm for computing option prices that we use requires that $\log G(\lambda)$ be linear in λ . Define $g(\lambda) = \log G(\lambda)$. For a given λ^* , note that for λ close to λ^* ,

$$g(\lambda) \simeq g(\lambda^*) + (\lambda - \lambda^*)g'(\lambda^*). \quad (\text{C.1})$$

Moreover,

$$\begin{aligned} g'(\lambda^*) &= \frac{G'(\lambda^*)}{G(\lambda^*)} \\ &= \frac{1}{G(\lambda^*)} \int_0^\infty b_\phi(\tau) \exp \{a_\phi(\tau) + b_\phi(\tau)\lambda^*\} d\tau. \end{aligned} \quad (\text{C.2})$$

The expression (C.2) has an interpretation: it is a weighted average of the coefficients $b_\phi(\tau)$, where the average is over τ , and the weights are proportional to $\exp \{a_\phi(\tau) + b_\phi(\tau)\lambda^*\}$. With

this in mind, we define the notation

$$b_\phi^* = \frac{1}{G(\lambda^*)} \int_0^\infty b_\phi(\tau) \exp \{a_\phi(\tau) + b_\phi(\tau)\lambda^*\} d\tau \quad (\text{C.3})$$

and the log-linear function

$$\hat{G}(\lambda) = G(\lambda^*) \exp \{b_\phi^*(\lambda - \lambda^*)\}. \quad (\text{C.4})$$

It follows from exponentiating both sides of (C.1) that

$$G(\lambda) \simeq \hat{G}(\lambda).$$

This log-linearization method differs from the more widely-used method of Campbell (2003), applied in continuous time by Chacko and Viceira (2005). However, in this application it is more accurate over the relevant range. This is not surprising, since we are able to exploit the fact that the true solution for the price-dividend ratio is known. In dynamic models with the EIS not equal to one, the solution is typically unknown.

We have computed implied volatilities using the loglinear approximation described above and by solving the expectation in (13) directly, by averaging over simulated sample paths. To keep the computation tractable, we assume a single jump size of -30%. The implied volatilities are indistinguishable.

C.2 Transform analysis

The normalized put option price is given as

$$P^n(\lambda_t, T - t; K^n) = E_t \left[\frac{\pi_T}{\pi_t} \left(K^n - \frac{S_T}{S_t} \right)^+ \right] \quad (\text{C.5})$$

It follows from $S_t = D_t G(\lambda_t)$, (9), (B.11), and (C.4) that

$$\begin{aligned}\frac{\pi_T}{\pi_t} &= \exp \left\{ - \int_t^T (\beta b \lambda_s - \eta) ds - \gamma \log \left(\frac{C_T}{C_t} \right) + b(\lambda_T - \lambda_t) \right\} \\ \frac{S_T}{S_t} &= \exp \left\{ \phi \log \left(\frac{C_T}{C_t} \right) + b_\phi^*(\lambda_T - \lambda_t) \right\},\end{aligned}$$

where η , b and b_ϕ^* are constants defined by (B.10), (2.2) and (C.3), respectively.²⁸

Then (C.5) can be rewritten as

$$\begin{aligned}P^n(\lambda_t, T - t; K^n) &= E_t \left[e^{-\int_t^T (\beta b \lambda_s - \eta) ds - \gamma(\log C_T - \log C_t) + b(\lambda_T - \lambda_t)} K^n \mathbf{1}_{\left\{ \frac{S_T}{S_t} \leq K^n \right\}} \right] \\ &\quad - E_t \left[e^{-\int_t^T (\beta b \lambda_s - \eta) ds + (\phi - \gamma)(\log C_T - \log C_t) + (b + b_\phi^*)(\lambda_T - \lambda_t)} \mathbf{1}_{\left\{ \frac{S_T}{S_t} \leq K^n \right\}} \right].\end{aligned}\quad (\text{C.6})$$

Note that

$$\mathbf{1}_{\left\{ \frac{S_T}{S_t} \leq K^n \right\}} = \mathbf{1}_{\left\{ b_\phi^*(\lambda_T - \lambda_t) + \phi(\log C_T - \log C_t) \leq \log K^n \right\}}.$$

Equation (C.6) characterizes the put option in terms of expectations that can be computed using the transform analysis of Duffie, Pan, and Singleton (2000). This analysis requires only the solution of a system of ordinary differential equations and a one-dimensional numerical integration. Below, we describe how we use their analysis.

To use the method of Duffie, Pan, and Singleton (2000), it is helpful to write down the following stochastic process, which, under our assumptions, is well-defined for a given λ_t .

$$X_\tau = \begin{bmatrix} \log C_{t+\tau} - \log C_t \\ \lambda_{t+\tau} \end{bmatrix}.$$

²⁸The equations for the CDR case are a special case of these. As is intuitive, $S_T/S_t = (C_T/C_t)^\phi$. The equation π_T/π_t is not as obvious because b and a themselves depend on σ_λ , but can be computed from (B.12). The CDR case as a limit of the SDR case is analyzed rigorously in Appendix B.2.

Note that the $\{X_\tau\}$ process is defined purely for mathematical convenience. Further define

$$\begin{aligned} R(X_\tau) &= X_\tau^\top [0, \beta b] - \eta = \beta b \lambda_{t+\tau} - \eta \\ d_1 &= [-\gamma, b]^\top \\ d_2 &= [\phi, b_\phi^*]^\top, \end{aligned}$$

and let

$$\mathcal{G}_{p,q}(y; X_0, T-t) = E \left[e^{-\int_0^{T-t} R(X_\tau) d\tau} e^{p^\top X_{T-t}} \mathbf{1}_{\{q^\top X_{T-t} \leq y\}} \right]. \quad (\text{C.7})$$

Note that $\{X_\tau\}$ is an affine process in the sense defined by Duffie et al. It follows that

$$\begin{aligned} P^n(\lambda, T-t; K^n) &= e^{-b\lambda} K^n E \left[e^{-\int_0^{T-t} R(X_\tau) d\tau + d_1^\top X_{T-t}} \mathbf{1}_{d_2^\top X_{T-t} \leq \log K^n + b_\phi^* \lambda} \right] \\ &\quad - e^{-(b+b_\phi^*)\lambda} E \left[e^{-\int_0^{T-t} R(X_\tau) d\tau + (d_1+d_2)^\top X_{T-t}} \mathbf{1}_{d_2^\top X_{T-t} \leq \log K^n + b_\phi^* \lambda} \right], \end{aligned}$$

and therefore

$$\begin{aligned} P^n(\lambda, T-t; K^n) &= e^{-b\lambda} \left(K^n \mathcal{G}_{d_1, d_2}(\log K^n + b_\phi^* \lambda, X_0, T-t) \right. \\ &\quad \left. - e^{-b_\phi^* \lambda} \mathcal{G}_{d_1+d_2, d_2}(\log K^n + b_\phi^* \lambda, X_0, T-t) \right), \end{aligned}$$

where $X_0 = [0, \lambda]$. The terms written using the function \mathcal{G} can then be computed tractably using the transform analysis of Duffie et al.

D The Poisson-normal model in continuous time

Backus, Chernov, and Martin (2011) specify a discrete-time model for consumption growth. This specification can be mapped naturally to continuous time. In their paper, the log consumption growth has a two-component structure. The first component is normal and the

second component is a Poisson mixture of normals. Consider the following continuous-time specification:

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

$$Z_t \sim N(\theta, \delta^2), \quad N_t \sim \text{Poisson}(\bar{\lambda})$$

Ito's Lemma implies

$$d \log \left(\frac{C_t}{C_0} \right) = \left(\mu - \frac{1}{2} \sigma^2 \right) dt + \sigma dB_t + Z_t dN_t$$

and therefore,

$$\log \left(\frac{C_t}{C_0} \right) = \left(\mu - \frac{1}{2} \sigma^2 \right) t + \sigma B_t + \sum_{i=1}^{N_t} Z_i$$

We can observe that the log consumption growth has a normal component and a jump component. Because the jump sizes are independent, and because each Z_t has a normal distribution,

$$\left(\sum_{i=1}^{N_t} Z_i \mid N_t = j \right) = \sum_{i=1}^j Z_i \sim N(j\theta, j\delta^2),$$

where j has a Poisson distribution. This is exactly the endowment process assumed by Backus, Chernov, and Martin (2011).

References

- Abel, Andrew, 1999, Risk premia and term premia in general equilibrium, *Journal of Monetary Economics* 43, 3–33.
- Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters Implied by Equity Index Options, *The Journal of Finance* 66, 1969–2012.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long-run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866.
- Barro, Robert J., and Tao Jin, 2011, On the size distribution of macroeconomic disasters, *Econometrica* 79, 1567–1589.
- Barro, Robert J., and José F. Ursúa, 2008, Macroeconomic crises since 1870, *Brookings Papers on Economic Activity* no. 1, 255–350.
- Barro, Robert J., and José F. Ursúa, 2009, Stock-market crashes and depressions, NBER Working Paper No. 14760.
- Bates, David S., 2000, Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181–238.
- Bates, David S., 2008, The market for crash risk, *Journal of Economic Dynamics and Control* 32, 2291–2321.
- Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, 2011, Explaining asset pricing puzzles associated with the 1987 market crash, *Journal of Financial Economics* 101, 552 – 573.

- Berkman, Henk, Ben Jacobsen, and John B. Lee, 2011, Time-varying rare disaster risk and stock returns, *Journal of Financial Economics* 101, 313–332.
- Black, Fischer, and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637–654.
- Bollerslev, Tim, and Viktor Todorov, 2011, Tails, Fears, and Risk Premia, *The Journal of Finance* 66, 2165–2211.
- Breeden, Douglas T., 1979, An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265–296.
- Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premia: Evidence from futures options, *Journal of Finance* 62, 1453–1490.
- Brown, Stephen J., William N. Goetzmann, and Stephen A. Ross, 1995, Survival, *The Journal of Finance* 50, 853–873.
- Buraschi, Andrea, and Alexei Jiltsov, 2006, Model Uncertainty and Option Markets with Heterogeneous Beliefs, *The Journal of Finance* 61, 2841–2897.
- Campbell, John Y., 2003, Consumption-based asset pricing, in G. Constantinides, M. Harris, and R. Stulz, eds.: *Handbook of the Economics of Finance, vol. 1b* (Elsevier Science, North-Holland).
- Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

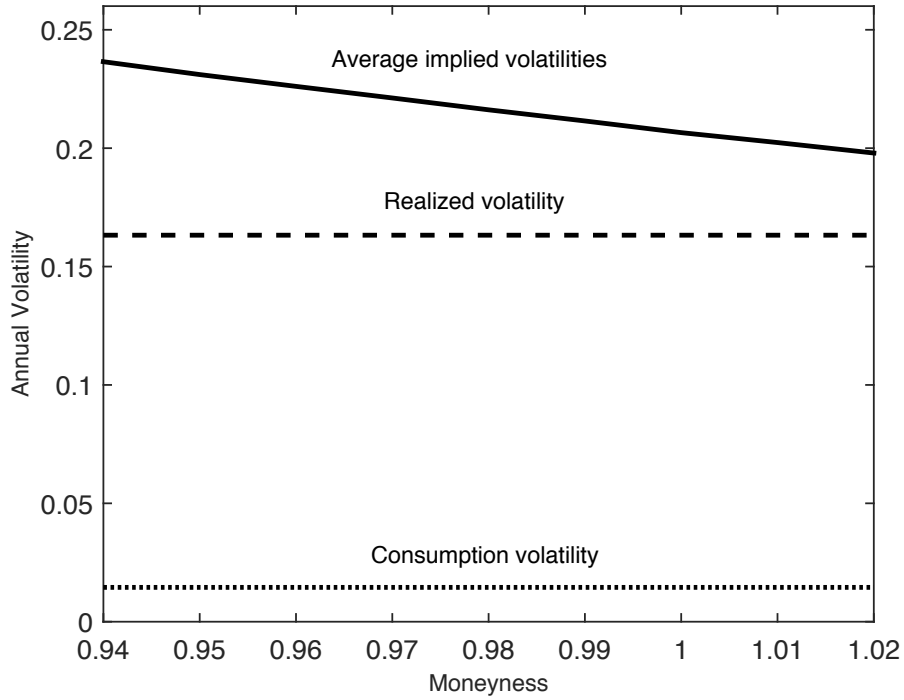
- Chacko, George, and Luis Viceira, 2005, Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets, *Review of Financial Studies* 18, 1369–1402.
- Chen, Hui, Winston Wei Dou, and Leonid Kogan, 2013, Measuring the “Dark Matter” in Asset Pricing Models, Working paper, MIT.
- Christoffersen, Peter, Steven Heston, and Kris Jacobs, 2009, The Shape and Term Structure of the Index Option Smirk: Why Multifactor Stochastic Volatility Models Work So Well, *Management Science* 55, 1914–1932.
- Christoffersen, Peter, and Kris Jacobs, 2004, The importance of the loss function in option valuation, *Journal of Financial Economics* 72, 291–318.
- Constantinides, George M., Jens Carsten Jackwerth, and Stylianos Perrakis, 2009, Mispricing of S&P 500 Index Options, *Review of Financial Studies* 22, 1247–1277.
- Coval, Joshua D., and Tyler Shumway, 2001, Expected option returns, *The Journal of Finance* 56, 983–1009.
- Cox, John C., Jonathan C. Ingersoll, and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–408.
- Drechsler, Itamar, 2013, Uncertainty, Time-Varying Fear, and Asset Prices, *The Journal of Finance* 68, 1843–1889.
- Drechsler, Itamar, and Amir Yaron, 2011, What’s vol got to do with it, *Review of Financial Studies* 24, 1–45.
- Du, Du, 2011, General equilibrium pricing of options with habit formation and event risks, *Journal of Financial Economics* 99, 400–426.

- Duffie, Darrell, and Larry G Epstein, 1992, Asset pricing with stochastic differential utility, *Review of Financial Studies* 5, 411–436.
- Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.
- Duffie, Darrell, and Costis Skiadas, 1994, Continuous-time asset pricing: A utility gradient approach, *Journal of Mathematical Economics* 23, 107–132.
- Dumas, Bernard, Jeff Fleming, and Robert E. Whaley, 1998, Implied Volatility Functions: Empirical Tests, *The Journal of Finance* 53, 2059–2106.
- Epstein, Larry, and Stan Zin, 1989, Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.
- Eraker, Bjørn, 2004, Do stock prices and volatility jump? Reconciling evidence from spot and option prices, *The Journal of Finance* 59, 1367–1404.
- Gabaix, Xavier, 2012, An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.
- Gao, George P., and Zhaogang Song, 2013, Rare disaster concerns everywhere, Working paper, Cornell University.
- Jorion, Philippe, and William N. Goetzmann, 1999, Global Stock Markets in the Twentieth Century, *The Journal of Finance* 54, 953–980.
- Kelly, Bryan T., Lubos Pastor, and Pietro Veronesi, 2014, The Price of Political Uncertainty: Theory and Evidence from the Option Market, Working paper, University of Chicago.

- Lettau, Martin, Sydney C. Ludvigson, and Jessica A. Wachter, 2008, The declining equity premium: What role does macroeconomic risk play?, *Review of Financial Studies* 21, 1653–1687.
- Liu, Jun, Jun Pan, and Tan Wang, 2005, An equilibrium model of rare-event premia and its implication for option smirks, *Review of Financial Studies* 18, 131–164.
- Longstaff, Francis A., and Monika Piazzesi, 2004, Corporate earnings and the equity premium, *Journal of Financial Economics* 74, 401–421.
- Manela, Asaf, and Alan Moreira, 2015, News implied volatility and disaster concerns, forthcoming, *Journal of Financial Economics*.
- Mehra, Rajnish, and Edward Prescott, 1985, The equity premium puzzle, *Journal of Monetary Economics* 15, 145–161.
- Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding predictability, *Journal of Political Economy* 112, 1–47.
- Naik, Vasanttilak, and Moon Lee, 1990, General equilibrium pricing of options on the market portfolio with discontinuous returns, *Review of Financial Studies* 3, 493–521.
- Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa, 2013, Crises and recoveries in an empirical model of consumption disasters, *American Economic Journal: Macroeconomics* 5, 35–74.
- Neuberger, Anthony, 2012, Realized Skewness, *Review of Financial Studies* 25, 3423–3455.
- Nowotny, Michael, 2011, Disaster begets crisis: The role of contagion in financial markets, Working paper, Boston University.

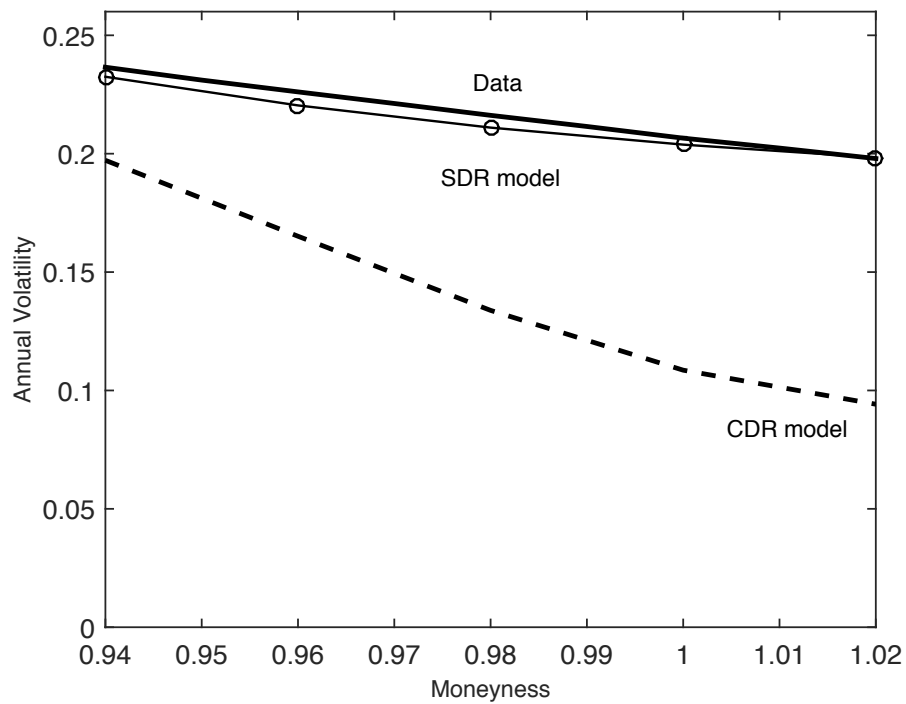
- Pan, Jun, 2002, The jump-risk premia implicit in options: evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.
- Rubinstein, Mark, 1994, Implied Binomial Trees, *The Journal of Finance* 49, 771–818.
- Santa-Clara, Pedro, and Shu Yan, 2010, Crashes, Volatility, and the Equity Premium: Lessons from S&P 500 Options, *Review of Economics and Statistics* 92, 435–451.
- Seo, Sang Byung, and Jessica A. Wachter, 2016, Do rare events explain CDX tranche spreads?, Working paper, University of Houston and University of Pennsylvania.
- Shaliastovich, Ivan, 2015, Learning, confidence, and option prices, *Journal of Econometrics* 187, 18–42.
- Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421–436.
- Tauchen, George, 2005, Stochastic volatility in general equilibrium, Working paper, Duke University.
- Tsai, Jerry, and Jessica A. Wachter, 2015, Disaster Risk and Its Implications for Asset Pricing, *Annual Review of Financial Economics* 7, 219–252.
- Wachter, Jessica A., 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *The Journal of Finance* 68, 987–1035.
- Weil, Philippe, 1990, Nonexpected utility in macroeconomics, *Quarterly Journal of Economics* 105, 29–42.
- Yan, Shu, 2011, Jump risk, stock returns, and slope of implied volatility smile, *Journal of Financial Economics* 99, 216–233.

Figure 1: Realized and implied volatility



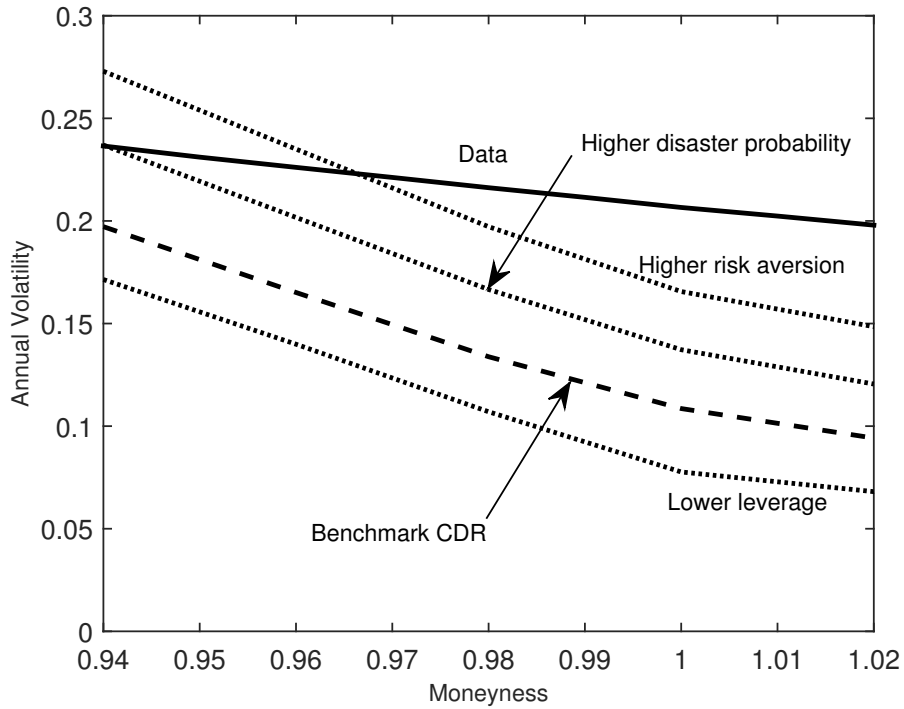
Notes: Implied volatilities are derived using the Black-Scholes formula from 3-month equity index options, and then averaged. Realized volatility is the standard deviation of annual log price changes for the S&P 500 index. Consumption volatility is the standard deviation of year-over-year changes in log aggregate real per-capita consumption of non-durables and services. Implied volatilities are shown as functions of moneyness, defined as the exercise price divided by the price of the index. Data are from 1996 to 2012.

Figure 2: Average implied volatilities in the SDR and CDR models



Notes: Average implied volatilities for 3-month options for the stochastic disaster risk (SDR) model, for the constant disaster risk (CDR) model and in the data. Average implied volatilities are shown as functions of moneyness, defined as the exercise price divided by the asset price.

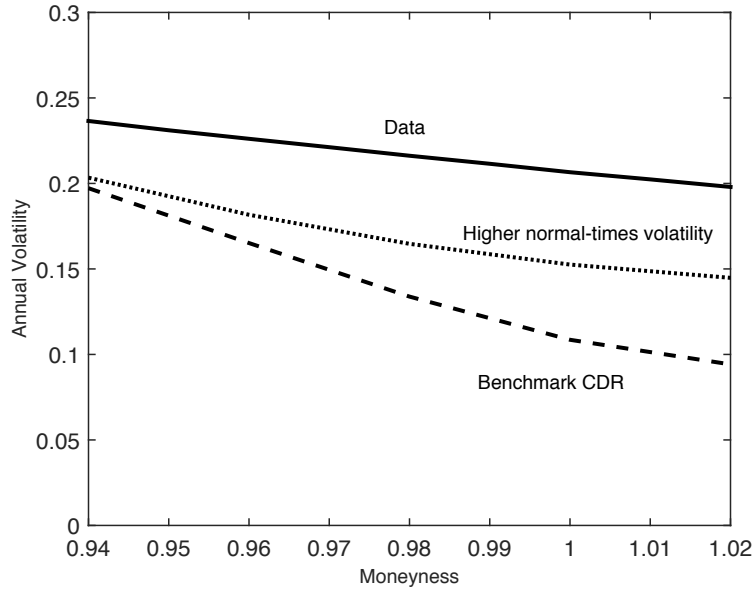
Figure 3: Comparative statics for the constant disaster risk (CDR) model




Notes: Average implied volatilities in the data (solid line), in the benchmark CDR model (dashed line), and in three alternative calibrations (dotted lines). In the first, risk aversion γ is equal to 4 rather than 3 (“higher risk aversion”). In the second, the disaster probability λ is equal to 5.6% rather than 3.6% (“higher disaster probability”). In the third, leverage ϕ is equal to 1.6 rather than 2.6 (“lower leverage”). All other parameters are kept at their benchmark level. Implied volatilities are shown as function of moneyness (exercise price over asset price).

Figure 4: The effect of greater normal-times consumption risk

Panel A: Implied volatilities

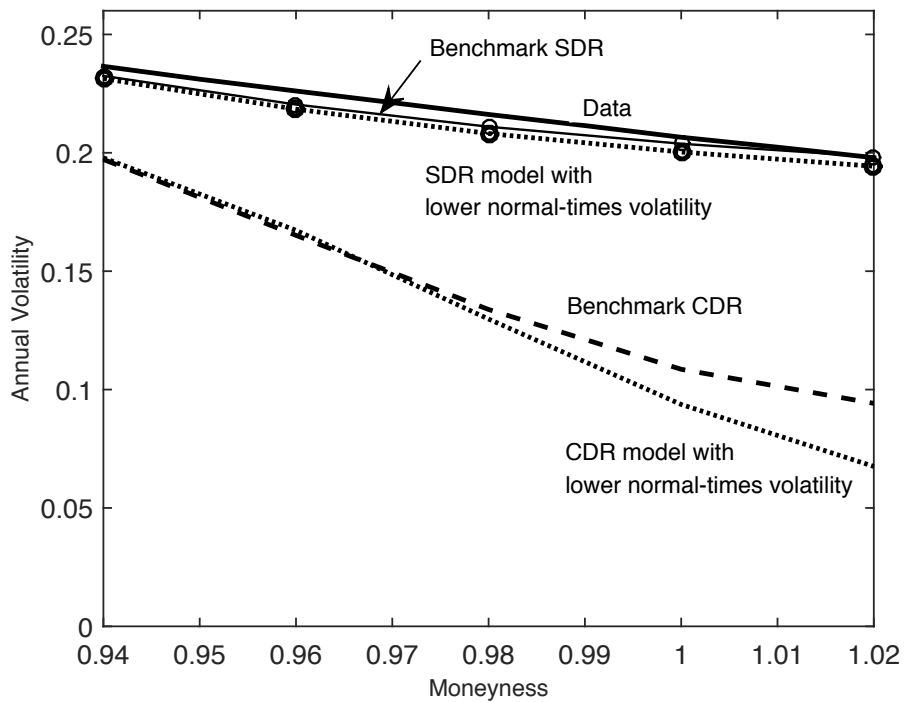


Panel B: Sampling distribution of consumption volatility



Notes: Panel A shows average implied volatilities for 3-month options as a function of moneyness (exercise price over asset price) for the benchmark constant disaster risk (CDR) model, for the CDR model with normal-times consumption volatility σ equal to 4% rather than 2% (“higher normal-times volatility”), and in the data. Panel B shows the histogram of volatility of log consumption growth from 10,000 simulations of the consumption process from the CDR model with $\sigma = 4\%$ with length designed to match the postwar data. The red dashed line is at 2%. The minimum volatility out of 10,000 distributions is 2.75%.

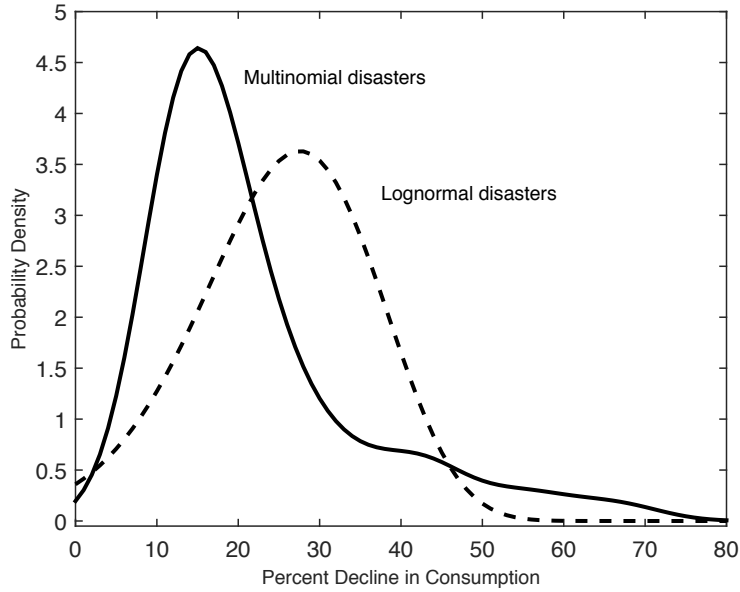
Figure 5: The effect of lower normal-times consumption risk




Notes: Average implied volatilities for 3-month options as a function of moneyness in the data (solid line), in the benchmark calibrations of the SDR and CDR models (solid line with circles and dashed line, respectively), and in a calibration of both models with normal-times consumption volatility σ set at 1% (dotted lines).

Figure 6: The role of the disaster distribution

Panel A: Disaster distribution

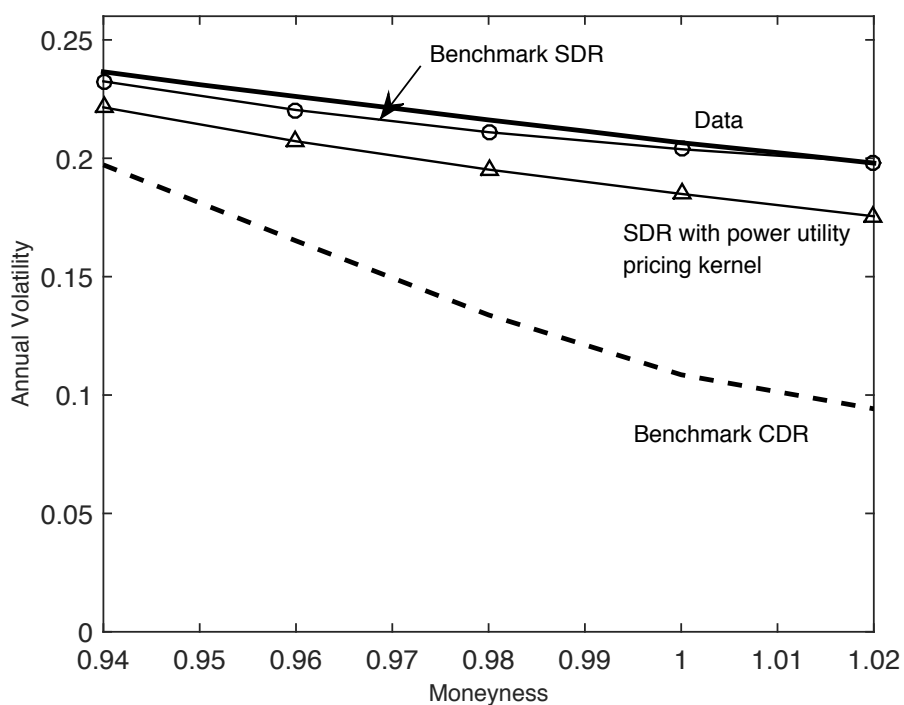


Panel B: Implied volatilities



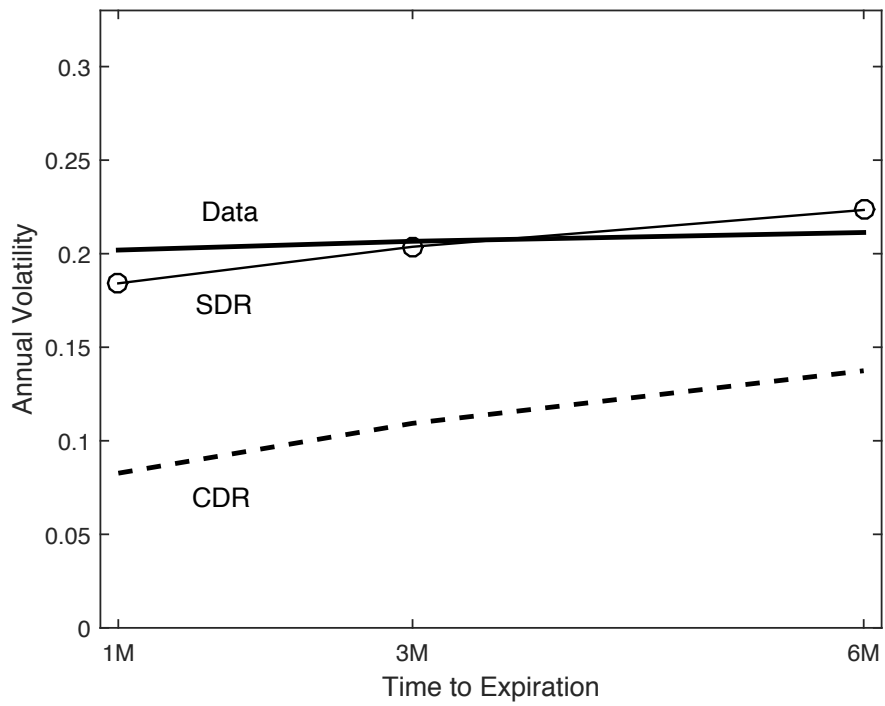
Notes: Panel A shows the density of percent consumption declines for the disaster distribution estimated by Barro and Ursúa (2008) and for a lognormal distribution with mean parameter 0.30 and standard deviation parameter 0.15. Panel B shows average implied volatilities for 3-month options as a function of moneyness in the data, in CDR and SDR models assuming the multinomial disaster distribution (the benchmark), and assuming the lognormal distribution.

Figure 7: The effect of the preference for early resolution of uncertainty.



Notes: Average implied volatilities for 3-month options as a function of moneyness computed for an SDR model where the pricing kernel is generated by time-additive (power) utility. The benchmark SDR model assumes an agent with a preference for an early resolution of uncertainty. In both models, the stock price is the same, and assumes a preference for an early resolution of uncertainty. Also shown are average implied volatilities in the CDR model and in the data.

Figure 8: The volatility term structure



Notes: Implied volatilities in the data, the SDR, and the CDR model for at-the-money (ATM) options as a function of time-to-expiration in months.

Table 1: Parameter values

Relative risk aversion γ	3.0
EIS ψ	1
Rate of time preference β	1.2%
Average growth in consumption (normal times) μ	2.52%
Volatility of consumption growth (normal times) σ	2%
Leverage ϕ	2.6
Average probability of a rare disaster $\bar{\lambda}$	3.55%
Mean reversion κ	0.080
Volatility parameter σ_λ	6.7%

Notes: Parameter values for the benchmark SDR model, in annual terms. The parameters for the benchmark CDR model are the same, except that σ_λ is equal to zero and κ is undefined.