Option prices in a model with stochastic disaster risk *

Sang Byung Seo
University of Pennsylvania

Jessica A. Wachter
University of Pennsylvania
and NBER

October 18, 2013

Abstract

Large rare shocks to aggregate consumption, namely, disasters, have been proposed as an explanation for the equity premium. However, recent work suggests that the consumption distribution required by this mechanism is inconsistent with the average implied volatility curve derived from option prices. We show that this apparent inconsistency can be resolved in a model with stochastic disaster risk. That is, we show that a model with a stochastic probability of disaster can explain average implied volatilities, despite being calibrated to consumption and aggregate market data alone. We also extend the stochastic disaster risk model to one that allows for variation in the risk of disaster at different time scales. We show that this extension allows the model to match variation in the level and slope of implied volatilities, as well as the average implied volatility curve.

*First draft: January 18, 2013. Seo: sangseo@wharton.upenn.edu; Wachter: jwachter@wharton.upenn.edu. We thank Hui Chen, Domenico Cuoco, Mikhail Chernov, Xavier Gabaix, Ivan Shaliastovich, Viktor Todorov and seminar participants at the 2013 NBER Summer Institute and at the Wharton school for helpful comments.

Option prices in a model with stochastic disaster risk

Large rare shocks to aggregate consumption, namely, disasters, have been proposed as an explanation for the equity premium. However, recent work suggests that the consumption distribution required by this mechanism is inconsistent with the average implied volatility curve derived from option prices. We show that this apparent inconsistency can be resolved in a model with stochastic disaster risk. That is, we show that a model with a stochastic probability of disaster can explain average implied volatilities, despite being calibrated to consumption and aggregate market data alone. We also extend the stochastic disaster risk model to one that allows for variation in the risk of disaster at different time scales. We show that this extension allows the model to match variation in the level and slope of implied volatilities, as well as the average implied volatility curve.

1 Introduction

Century-long evidence indicates an economically significant equity premium, namely, the expected return from holding equities over short-term debt.[1] The source of this equity premium has been a subject of debate for nearly thirty years.[2] One place to look for such a source is in options data. By holding equity and a put option, an investor can, at least in theory, eliminate the downside risk in equities. For this reason, it is appealing to explain both options data and standard equity returns together with a single model.

Such an approach is arguably of particular importance for a class of models that explain the equity premium through the mechanism of consumption disasters (e.g. Rietz (1988), Barro (2006) and Weitzman (2007)). In these models, consumption growth rates and thus equity returns are subject to shocks that are rare and large. Options (assuming away, for the moment, the potentially important question of counterparty risk), offer a way to hold equities while eliminating the exposure to disasters. Thus it is of interest to know whether these models have the potential to explain option prices as well as equity prices.

Option prices are also of interest in their own right because their convex payoff structure gives more information on the distribution of returns than equities alone. While international consumption data used by Barro (2006) and others gives some information on the probabilities of rare events that U.S. market participants impute into prices, option prices can potentially offer a second, largely independent source of information. Indeed, a large and sophisticated literature concludes that large sudden moves in stock returns are important for explaining option prices.[3] In the options literature, these large sudden changes in price are

[1] See Mehra and Prescott (1985). Siegel (1994) presents historical evidence on equity performance; Campbell (2003) presents international evidence.

[2] Surveys include Mehra and Prescott (2003) and Kocherlakota (1996).

[3] One branch of this literature focuses on reduced-form models that link stock returns to option prices (Bakshi, Cao, and Chen (1997), Bates (2000), At-Sahalia, Wang, and Yared (2001), Pan (2002), Broadie,

called jumps rather than disasters. However, in both literatures, what is being referred to are events whose size and probability of occurrence renders a normal distribution essentially impossible.

Motivated by the parallels between these very different literatures, Backus, Chernov, and Martin (2011) study option prices in a rare events model similar to that of Rietz (1988) and Barro (2006). They find, however, that the resulting options prices are far from their data counterparts. They argue that options data appear to be inconsistent with the hypothesis of large and rare shocks to equity returns. In particular, the implied volatilities resulting from rare events models are lower than in the data, and are far more downward sloping as a function of the strike price. Further, they show through the lens of their model, that events of the magnitude required to explain the equity premium are virtually ruled out by option prices. These results suggest that rare events are not the source of risk behind the equity premium.

Like Barro (2006) and Rietz (1988), Backus, Chernov, and Martin (2011) assume that the probability of a disaster occurring is constant. Such a model can explain the equity premium, but cannot account for other features of equity markets, such as the volatility. Recent rare events models therefore introduce dynamics that can account for equity volatility (Gabaix (2012), Gourio (2011), Wachter (2013)).[4] We derive option prices in a model based on that of Wachter (2013), and in a significant generalization of this model that allows for variation in the risk of disaster at different time scales. We show that allowing for a stochastic probability

Chernov, and Johannes (2007), Santa-Clara and Yan (2010)). Another branch makes use of the information in very short-term options to separate out the jump component of stock returns (Carr and Wu (2003), Bollerslev and Todorov (2011)).

[4]These models assume an exogenous process for time-variation in the probability of a disaster. However, time-variation in the probability can also be driven by learning (Veronesi (2004)) and by heterogeneity in beliefs (Chen, Joslin, and Tran (2012)).

of disaster has dramatic effects on implied volatilities. Namely, rather than being much lower than in the data, the implied volatilities are at about the same level. The slope of the implied volatility curve, rather than being far too great, also matches that of the data. We then apply the model to understanding other features of option prices in the data, such as time-variation in the level and slope of the implied volatility curve. We show that the model can account for these features of the data as well.

We use the fact that our model generalizes the previous literature to better understand the difference in results. Because of time-variation in the disaster probability, the model endogenously produces the stock price changes that occur during normal times and that are reflected in option prices. These changes are absent in iid rare event models, because, during normal times, the volatility of stock returns is equal to the (low) volatility of dividend growth. Moreover, by assuming recursive utility, the model implies a premium for assets that covary negatively with volatility. This makes implied volatilities higher than what they would be otherwise.

Our findings relate to those of Gabaix (2012), who also reports average implied volatilities in a model with rare events.[5] Our model is conceptually different in that we assume recursive utility and time-variation in the probability of a disaster. Gabaix assumes a linearity-generating process (Gabaix (2008)) with a power utility investor. In his calibration, the sensitivity of dividends to changes in consumption is varying; however, the probability of a disaster is not. There are several important implications of our choice of approach. First, the tractability of our framework implies that we can use the same model for pricing options as we do for equities. Moreover, our model nests the simpler one with constant risk of disaster, allowing us to uncover the reason for the large difference in option prices between the models. We also show that the assumption of recursive utility is important for the ac-

[5]Recent work by Nowotny (2011) reports average implied volatilities as well. Nowotny focuses on the implications of self-exciting processes for equity markets rather than on option prices.

counting for the average implied volatility curve. Finally, we go beyond the average implied volatility curve for 3-month options, considering the time series in the level and slope of the implied volatility curve across different maturities. These considerations lead us naturally to a two-factor model for the disaster probability.

Other recent work explores implications for option pricing in dynamic endowment economies. Benzoni, Collin-Dufresne, and Goldstein (2011) derive options prices in a Bansal and Yaron (2004) economy with jumps to the mean and volatility of dividends and consumption. The jump probability can take on two states which are not observable to the agent. Their focus is on the learning dynamics of the states, and the change in option prices before and after the 1987 crash, rather than on matching the shape of the implied volatility curve. Du (2011) examines options prices in a model with external habit formation preferences as in Campbell and Cochrane (1999) in which the endowment is subject to rare disasters that occur with a constant probability. His results illustrate the difficulty of matching implied volatilities assuming either only external habit formation, or a constant probability of a disaster. Also related is the work of Drechsler and Yaron (2011), who focus on the volatility premium (see Bollerslev, Tauchen, and Zhou (2009)) and its predictive properties. In the work of Drechsler and Yaron, the equity premium arises from the combined effect of persistent growth rates in consumption, high risk aversion, and a high elasticity of intertemporal substitution (long-run risk), as opposed to rare events.

A related strand of research on endowment economies focuses on uncertainty aversion or exogenous changes in confidence. Drechsler (2012) builds on the work of Bansal and Yaron (2004), but incorporates dynamic uncertainty aversion. He argues that uncertainty aversion is important for matching implied volatilities. Shaliastovich (2009) shows that jumps in confidence can explain option prices when investors are biased toward recency. These papers build on earlier work by Bates (2008) and Liu, Pan, and Wang (2005), who conclude that it

is necessary to introduce a separate aversion to crashes to simultaneously account for data on options and on equities. Buraschi and Jiltsov (2006) explain the pattern in implied volatilities using heterogenous beliefs. Unlike these papers, we assume a rational expectations investor with standard (recursive) preferences. The ability of the model to explain implied volatilities arises from time-variation in the probability of a disaster rather than a premium associated with uncertainty.

The remainder of this paper is organized as follows. Section 2 discusses our model for disaster risk and equity prices, which we call the stochastic disaster risk (SDR) model. Section 3 compares the implications of the SDR model with the implications of a constant disaster risk (CDR) model for average implied volatilities. Section 4 extends the SDR model to a model in which two factors at different time scales drive the risk of a disaster. We show that this two-factor SDR model can fit qualitative features of the option data that the SDR model cannot. Section 5 concludes.

2 A stochastic disaster risk model

2.1 Assumptions

In this section we describe a model with stochastic disaster risk (SDR). We assume an endowment economy with complete markets and an infinitely-lived representative agent. Aggregate consumption (the endowment) solves the following stochastic differential equation

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

where B_t is a standard Brownian motion and N_t is a Poisson process with time-varying intensity λ_t. This intensity follows the process

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

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 of a jump. We thus will use the terminology probability and intensity interchangeably, while keeping in mind the that the relation is an approximate one.

The size of a jump, provided that a jump occurs, is determined by Z_t. We assume Z_t is a random variable whose time-invariant distribution ν is independent of N_t, B_t and $B_{\lambda,t}$. We will use the notation E_ν to denote expectations of functions of Z_t taken with respect to the ν-distribution. The t subscript on Z_t will be omitted when not essential for clarity.

We will assume a recursive generalization of power utility that allows for preferences over the timing of the resolution of uncertainty. Our formulation comes from Duffie and Epstein (1992), and we consider a special case in which the parameter that is often interpreted as the elasticity of intertemporal substitution (EIS) is equal to 1. That is, we define continuation utility V_t for the representative agent using the following recursion:

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

where

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

The parameter β is the rate of time preference. We follow common practice in interpreting γ as relative risk aversion. 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 Solving for asset prices

We will solve for asset prices using the state-price density, π_t.[6] Duffie and Skiadas (1994) characterize the state-price density as

$$\pi_t = \exp\left\{\int_0^t \frac{\partial}{\partial V} f\left(C_s, V_s\right) ds\right\} \frac{\partial}{\partial C} f\left(C_t, V_t\right). \tag{5}$$

There is an equilibrium relation between utility V_t, consumption C_t and the disaster probability λ_t. Namely,

$$V_t = \frac{C_t^{1-\gamma}}{1-\gamma} e^{a+b\lambda_t},$$

where a and b are constants given by

$$a = \frac{1-\gamma}{\beta}\left(\mu - \frac{1}{2}\gamma\sigma^2\right) + b\frac{\kappa\bar{\lambda}}{\beta} \tag{6}$$

$$b = \frac{\kappa+\beta}{\sigma_\lambda^2} - \sqrt{\left(\frac{\kappa+\beta}{\sigma_\lambda^2}\right)^2 - 2\frac{E_\nu\left[e^{(1-\gamma)Z}-1\right]}{\sigma_\lambda^2}}. \tag{7}$$

It follows that

$$\pi_t = \exp\left(\eta t - \beta b \int_0^t \lambda_s ds\right) \beta C_t^{-\gamma} e^{a+b\lambda_t}, \tag{8}$$

where $\eta = -\beta(a+1)$. Details are provided in Appendix A.1.

Following Backus, Chernov, and Martin (2011), we assume a simple relation between dividends and consumption: $D_t = C_t^\phi$, for leverage parameter ϕ.[7] Let $F(D_t, \lambda_t)$ be the value of the aggregate market (it will be apparent in what follows that F is a function of D_t and λ_t). It follows from no-arbitrage that

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

[6] Other work on solving for equilibria in continuous-time models with recursive utility includes Benzoni, Collin-Dufresne, and Goldstein (2011), Eraker and Shaliastovich (2008), Fisher and Gilles (1999) and Schroder and Skiadas (1999).

[7] This implies that dividends respond more than consumption to disasters, an assumption that is plausible given the U.S. data (Longstaff and Piazzesi (2004)).

The stock price can be written explicitly as

$$F(D_t, \lambda_t) = D_t G(\lambda_t), \tag{9}$$

where the price-dividend ratio G is given by

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

for functions $a_\phi(\tau)$ and $b_\phi(\tau)$ given by:

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

where

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

(see Wachter (2013)). We will often use the abbreviation $F_t = F(D_t, \lambda_t)$ to denote the value of the stock market index at time t.

2.3 Implied volatilities in the stochastic disaster risk model

Let $P(F_t, \lambda_t, \tau; K)$ denote the time-t price of a European put option on the stock market index with strike price K and expiration $t+\tau$. For simplicity, we will abbreviate the formula for the price of the dividend claim as $F_t = F(D_t, \lambda_t)$. Because the payoff on this option at expiration is $(K - F_{t+\tau})^+$, it follows from the absence of arbitrage that

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

Let $K^n = K/F_t$, the normalized strike price (or "moneyness"), and define

$$P^n(\lambda_t, T-t; K^n) = E_t\left[\frac{\pi_T}{\pi_t}\left(K^n - \frac{F_T}{F_t}\right)^+\right]. \tag{10}$$

We will establish below that P^n is indeed a function of λ_t, time to expiration and moneyness alone Clearly $P_t^n = P_t/F_t$. Because our ultimate interest is in implied volatilities, and because, in the formula of Black and Scholes (1973), normalized option prices are functions of the normalized strike price (and the volatility, interest rate and time to maturity), it suffices to calculate P_t^n.[8]

Returning to the formula for P_t^n, we note that, from (9), it follows that

$$\frac{F_T}{F_t} = \frac{D_T}{D_t}\frac{G(\lambda_T)}{G(\lambda_t)}. \tag{11}$$

Moreover, it follows from (8) that

$$\frac{\pi_T}{\pi_t} = \left(\frac{C_T}{C_t}\right)^{-\gamma} \exp\left\{\int_t^T (\eta - \beta b \lambda_s)\, ds + b(\lambda_T - \lambda_t)\right\}. \tag{12}$$

At time t, λ_t is sufficient to determine the distributions of consumption and dividend growth between t and T, as well as the distribution of λ_s for $s = t, \ldots, T$. It follows that normalized put prices (and therefore implied volatilities) are a function of λ_t, the time to expiration, and moneyness.

[8]Given stock price F, strike price K, time to maturity $T-t$, interest rate r, and dividend yield y, the Black-Scholes put price is defined as

$$\mathrm{BSP}(F, K, T-t, r, y, \sigma) = e^{-r(T-t)} K N(-d_2) - e^{-y(T-t)} F N(-d_1)$$

where

$$d_1 = \frac{\log(F/K) + \left(r - y + \sigma^2/2\right)(T-t)}{\sigma\sqrt{T-t}} \quad \text{and} \quad d_2 = d_1 - \sigma\sqrt{T-t}$$

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

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

where r_t^b is the model's analogue of the Treasury Bill rate, which allows for a probability of a default in case of a disaster (see Barro (2006); as in that paper we assume a default rate of 0.4).

Appendix B describes the calculation of (10). We first approximate the price-dividend ratio $G(\lambda_t)$ by a log-linear function of λ_t, As the Appendix describes, this approximation is highly accurate. We can then apply the transform analysis of Duffie, Pan, and Singleton (2000) to calculate put prices.

The implied volatility curve in the data represents an average of implied volatilities at different points in time. We follow the same procedure in the model, calculating an unconditional average implied volatility curve. To do so, we first solve for the implied volatility as a function of λ_t. We numerically integrate this function over the stationary distribution of λ_t. This 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)).

2.4 The constant disaster risk model

Taking limits in the above model as σ_λ approaches zero implies a model with a constant probability of disaster (Appendix A.2 shows that this limit is indeed well-defined and is what would be computed if one were to solve the constant disaster risk model from first principles). We use this model to evaluate the role that stochastic disaster risk plays in the model's ability to match the implied volatility data. We refer to this model in what follows as the CDR (constant disaster risk) model, to distinguish it from the more general SDR model.

The CDR model is particularly useful in reconciling our results with those of Backus, Chernov, and Martin (2011). Backus et al. solve a model with a constant probability of a jump in consumption, calibrated in a manner similar to Barro (2006). They call this the consumption-based model, to distinguish it from the reduced-form options-based model which is calibrated to fit options data. While Backus et al. assume power utility, their model can be rewritten as one with recursive utility with an EIS of one. The reason is that

the endowment process is iid. In this special case, the EIS and the discount rate are not separately identified. Specifically, a power utility model has an identical stochastic discount factor, and therefore identical asset prices to a recursive utility model with arbitrary EIS as long as one can adjust the discount rate (Appendix A.3).[9]

3 Comparison with the data

3.1 Data

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. 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 (2009)).

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

[9] Backus et al. define option payoffs not in terms of the price, but in terms of the total return. Given the parameter range of interest, this distinction matters little for out-of-the money put options. For consistency with their study, we use their definition and pricing method when reporting results for the benchmark CDR calibration.

prices and times to expiration.[10] Specifically, at each date in the sample, we regress implied volatilities on a polynomial in strike price K and maturity T:

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

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.

3.2 Calibration

Table 1 shows the parameter values for the SDR and the benchmark CDR model. The parameters for the benchmark CDR model are as in the consumption-based model of Backus, Chernov, and Martin (2011), discussed above. This choice has two advantages. First, Backus et al. set the parameters of their model to match the equity premium and the variance of stock returns. Because the parameters in the SDR model also implies that the model can match the equity premium and stock market volatility (Wachter (2013)), the two models are in a sense equivalent. Second, this comparison will also allow us to reconcile our results with theirs.[11] Wn what follows, we will also compare the SDR model to a CDR model with the same parameters, except that σ_λ is set equal to zero.

As Table 1 shows, the two calibrations differ in their relative risk aversion, in the volatility of normal-times consumption growth, in leverage, in the probability of a disaster, and of course in whether the probability is time-varying. The net effect of some of these differences

[10] See Dumas, Fleming, and Whaley (1998), Christoffersen and Jacobs (2004) and Christoffersen, Heston, and Jacobs (2009).

[11] As described in the previous section, we consider a calibration that is isomorphic to that of Backus, Chernov, and Martin (2011) in which the EIS is equal to 1. This implies a discount rate of 0.0189, given that they calibrate their model assuming a riskfree rate of 2%. In contrast power utility implies a negative discount rate, reflecting the riskfree rate puzzle (Weil (1990)).

turns out to be less important than what one may think: for example, higher risk aversion and lower disaster probability roughly offset each other. We will explore the implications of leverage and volatility in what follows.

The two models also assume different disaster distributions. For the SDR model, the disaster distribution is multinomial, and taken from Barro and Ursua (2008) based on actual consumption declines. A histogram of this distribution is shown in Figure 1. the benchmark CDR model assumes that consumption declines are log-normal. For comparison, we plot the smoothed density for the SDR model along with the density of the consumption-based model in Figure 2. Compared with the lognormal model, the SDR model has more mass over small declines in the 10–20% region, and more mass over large declines in the 50-70% region.

3.3 Results

Figure 3 shows the resulting implied volatilities as a function of moneyness, as well as implied volatilities in the data. Confirming previous results, we find that the CDR model leads to implied volatilities that are dramatically different from those in the data. First, the implied volatilities are too low, even though the model was calibrated to match the volatility of equity returns. Second, they exhibit a strong downward slope as a function of the strike price. While there is a downward slope in the data, it is not nearly as large. As a result, implied volatilities for at-the-money (ATM) options in the CDR model are less than 10%, far below the option-based implied volatilities, which are over 20%.

In contrast, the SDR model can explain both ATM and OTM (out-of-the-money) implied volatilities. For OTM options (with moneyness equal to 0.94), the SDR model gives an implied volatility of 23%, close to the data value of 24%. There is a downward slope, just as in the data, but it is much smaller than that of the CDR model. ATM options have implied

volatilities of about 21% in both the model and the data.

Figure 3 shows population values for implied volatilities in the SDR model. This is appropriate in that the data averages implied volatility curves over the sample period. However, one might be concerned that, in the model, the population averages are skewed by periods during which disaster risk is relatively high. This would affect the robustness of the result. Figure 4 shows conditional implied volatilities in the SDR model for the median, the 20th, and the 80th percentile value. A higher disaster probability implies higher option prices and thus higher implied volatility. Most importantly, the curve at the median value also lies near to that in the data.

3.4 Discussion

There are a number of differences between this model and the CDR model. In this section, we discuss which of these differences is primarily responsible for the change in implied volatilities.

In their discussion, Backus, Chernov, and Martin (2011) emphasize the role of very bad consumption realizations as a reason for the poor performance of the disaster model. Therefore, this seems like an appropriate place to start. The disaster distribution in the SDR benchmark actually implies a slightly higher probability of extreme events than the benchmark CDR model (Figure 2). However, the benchmark CDR model has much higher leverage: the leverage parameter is 5.1 for the CDR calibration versus 2.6 for the SDR calibration. Leverage does not affect consumption but it affects dividends, and therefore stock and option prices. A higher leverage parameter implies that dividends will fall further in the event of a consumption disaster. It is reasonable, therefore, to attribute the difference in the implied volatilities to the difference in the leverage parameter.

Figure 5 tests this directly by showing option prices in the CDR model for leverage of 5.1 and for leverage of 2.6 (denoted "lower leverage") in the figure. Surprisingly, the slope

for the calibration with leverage of 2.6 is slightly higher than the slope for leverage of 5.1. Lowering leverage results in a downward shift in the level of the implied volatility curve, not the slope. Thus the difference in leverage cannot be the explanation for why the slope in our model is lower than the slope for CDR.

Why does the change in leverage result in a shift in the level of the curve? It turns out that in the CDR model, changing normal-times volatility has a large effect. Leverage affects both the disaster distribution and normal-times volatility. Lowering leverage has a large effect on normal-times volatility and thus at-the-money options. This is why the level of the curve is lower, and the slope is slightly steeper.[12]

To further consider the role of normal-times volatility, we explore the impact of changing the consumption volatility parameter σ. In the benchmark CDR comparison, consumption volatility is equal to the value of consumption volatility over the 1889–2009 sample, namely 3.5%. Most of this volatility is accounted for by the disaster distribution, because, while the disasters are rare, they are severe. Therefore normal-times volatility is 1%, lower than the U.S. consumption volatility over the post-war period. The SDR model is calibrated differently; following Barro (2006), the disaster distribution is determined based on international macroeconomic data, and the normal-times distribution is set to match postwar volatility in developed countries. The resulting normal-times volatility is 2%. To evaluate the effect of this difference, we solve for implied volatilities in the CDR model with leverage of 5.1 and normal-times volatility of 2%. In Figure 5, the result is shown in the line denoted "higher normal-times volatility."

As Figure 5 shows, increasing the normal-times volatility of consumption growth in the CDR model has a noticeable effect on implied volatilities: The implied volatility curve is

[12]Yan (2011) shows analytically that, as the time to expiration approaches zero, the implied volatility is equal to the normal-times volatility in the stock price, while the slope is inversely related to the normal-times volatility of the stock price.

higher and flatter. The change in the level reflects the greater overall volatility. The change in the slope reflects the greater probability of small, negative outcomes. However, the effect, while substantial, is not nearly large enough to explain the full difference. The level of the "higher normal-times volatility" smile is still too low and the slope is too high compared with the data.

While raising the volatility of consumption makes the CDR model look somewhat more like the SDR model (though it does not account for the full difference), it is not the case that lowering the volatility of consumption makes the SDR model more like the CDR model. Namely, reducing σ to 1% (which would imply a normal-times consumption volatility that is lower than in the post-war data) has almost no effect on the implied volatility curve of the SDR model. There are two reasons why this parameter affects implied volatilities differently in the two cases. First, the leverage parameter is much lower in the SDR model than the CDR model. Second, volatility in the SDR model comes from time-variation in discount rates (driven by λ_t) as well as in payouts ($\phi\sigma$). The first of these terms is much larger than the second.[13]

We also consider results from a CDR model that based on taking the limit of our SDR calibration as σ_λ approaches zero. The results are shown in Figure 6. The resulting implied volatility smile is also too low and steep as compared with the data. Lower leverage and higher normal-times volatility combine to make this slope similar to that of the benchmark CDR calibration, at least for out-of-the money options.

Given these results, it appears that the superior fit of the SDR model must arise from the fact that disaster risk is stochastic. A consequence of stochastic disaster risk is high stock

[13]To be precise, total 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.

market volatility, not just during occurrences of disasters, but during normal periods as well. This is reflected in the relative shallowness of the volatility smile: while the existence of disasters leads to an upward slope for out-of-the money put options, high normal-period volatility implies that the level is high for put options that are in the money or only slightly out of the money. The same mechanism, and indeed the same parameters that allow the model to match the level of realized stock returns enable the model to match implied volatilities.

The SDR model contains an additional mechanism that aids in the match to implied volatilities that is embedded in recursive utility. Recursive utility plays a number of roles in the model, including enabling the model to match realized volatilities; without recursive utility, the price-dividend ratio would not fall on an increased risk of rare disasters, because, at reasonable parameter values, the riskfree rate effect would be larger than the risk premium effect. However, there is a separate effect for options; As shown in Section 2.2, the state price density depends on the probability of disaster. Thus risk premia depend on covariances with this probability: assets that increase in price when the probability rises will be a hedge. Options are such an asset. Indeed, an increase in the probability of a rare disaster raises option prices, while at the same time increasing marginal utility. This can be seen in Figure 4; higher implied volatilities during disaster periods reflect higher option prices.

To directly assess the magnitude of this effect, we solve for option prices using the same process for the stock price and the dividend yield, but with a pricing kernel adjusted to set the above effect equal to zero. Risk premia in the model arise from covariances with the

pricing kernel. We replace the pricing kernel in (8) with one in which $b = 0$.[14] Because b determines the risk premium due to covariance with λ_t, setting $b = 0$ will shut off this effect. Indeed, as Figure 7 shows, setting $b = 0$ significantly reduces option prices, and hence implied volatilities.

Our results show that a model with stochastic disaster risk can fit implied volatilities, thereby addressing one issue raised by Backus, Chernov, and Martin (2011). Backus et al. raise a second issue as well. Assuming power utility and iid consumption growth, they back out a distribution for the left tail of consumption growth from option prices (we will call this the "option-implied consumption distribution").[15] Based on this distribution, they conclude that the probabilities of negative jumps to consumption are much larger, and the magnitudes much smaller, than implied by the international macroeconomic data used by Barro (2006) and Barro and Ursua (2008). In fact, events of the magnitude indicated by the international macroeconomic data are virtually ruled out.

The resolution of this second issue is clearly related to the first. For if a model (like the one we describe) can explain average implied volatilities while assuming a disaster distribution from macroeconomic data, then this macroeconomic disaster distribution is one possible consumption distribution that is consistent with the implied volatility curve. Namely, the inconsistency between the extreme consumption events in the macroeconomic data and option prices can be resolved by relaxing the iid assumption. Of course, it may be that the option-implied consumption growth distribution is also correct. This distribution is, after

[14]Note that a and η also depend on b: these expressions are also changed in the experiment. While it may first appear that b should also affect the riskfree rate, this does not occur in the model with EIS= 1. The riskfree rate satisfies a simple expression

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

[15]The basis for this analysis is the Merton (1976) model for returns, fitted to average implied volatilities.

all, consistent with option prices, 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, it turns out that this consumption distribution can be ruled out based on other data: because it assumes that negative consumption jumps are relatively frequent (as they must be to explain the equity premium), some would have occurred in the 60-year postwar period in the U.S. However, the unconditional volatility of consumption growth in the U.S. was less than 2%. Under the option-implied consumption growth distribution, there is less than a 1 in 100,000 chance of observing a 60-year period with volatility this low.[16]

We now return to the broader question of the mechanisms in our model that allow a match to implied volatilities. Previous work suggests that allowing for stochastic volatility (and time-varying moments more generally) does not appear to affect the shape of the implied volatility curve.[17] How is it, then, that this paper comes to such a different conclusion? The reason may arise from the fact that the previous literature mainly focused on reduced-form models, in which the jump dynamics and volatility of stock returns are freely chosen. However, in an equilibrium model like the present one, stock market volatility arises

[16] Another possible objection to this distribution is based on the iid assumption itself. Neuberger (2012) presents such an argument based the lack of decay in skewness as the measurement horizon grows. We discuss his results and their connection to ours in Section 4.3.3.

[17] 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."

endogenously from the interplay between consumption and dividend dynamics and agents' preferences. While it is possible to match the volatility of stock returns and consumption in an iid model, this can only be done (given the observed data) by having all of the volatility occur during disasters. In such a model it is not possible to generate sufficient stock market volatility in normal times to match either implied or realized volatilities. Thus, while in the reduced-form literature, the difference between iid and dynamic models principally affects the conditional moments of volatility, in the equilibrium literature, the difference affects the level of volatility itself.

4 Generalization

4.1 Why generalize the model?

The previous sections show how introducing time variation into conditional moments can substantially alter the implications of rare disasters for implied volatilities. The model presented there was parsimonious, with a single state variable following a square root process.

Closer examination of Figure 4 suggests an aspect of options data that may be difficult to fit to this model. This figure shows that implied volatilities for λ_t equal to the median and for the 20th and 80th percentile value. Implied volatilities increase almost in parallel as λ_t increases. That is, ATM options are affected by an increase in the rare disaster probability almost as much as out-of-the-money options. The model therefore implies that there should be little variation in the slope of the implied volatility curve.

Figure 8 shows the historical time series of implied volatilities computed on one month ATM and OTM options with moneyness of 0.85. Panel C shows the difference in the implied probabilities, a measure of the slope of the implied volatility curve. Defined in this way, the average slope is 12%, with a volatility of 2%. Moreover, the slope can rise as high as

18% and fall as low as 6%.[18] While the SDR model can explain the average slope, it seems unlikely that it would be able to account for the time-variation in the slope, at least under the current calibration. Moreover, comparing Panel C with Panels A and B of Figure 8 indicates that the slope varies independently of the level of implied volatilities. Thus it is unlikely that any model with a single state variable could account for these data.

The mechanism in the SDR model that causes time-variation in rare disaster probabilities is identical to the mechanism that leads to volatility in normal times. Namely, when λ_t is high, rare disasters are more likely *and* returns are more volatile. In order to account for the data, a model must somehow decouple the volatility of stock returns from the probability of rare events. This is challenging, because volatility endogenously depends on the probability of rare events. Indeed, the main motivation for assuming time-variation in the probability of rare events is to generate volatility in stock returns that seems otherwise puzzling. Developing such a model is our goal in this section of the paper.

4.2 A two-factor stochastic disaster risk model

In this section, we introduce a mechanism that decouples the volatility of stock returns from the probability of rare events. We assume the same stochastic process for consumption (1) and for dividends. We now assume, however, that the probability of a rare event follows the process

$$d\lambda_t = \kappa_\lambda(\xi_t - \lambda_t)dt + \sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t}, \tag{13}$$

where ξ_t solves the following stochastic differential equation

$$d\xi_t = \kappa_\xi(\bar{\xi} - \xi_t)dt + \sigma_\xi\sqrt{\xi_t}dB_{\xi,t}. \tag{14}$$

[18] Figure 4, for purposes of comparison with prior work, shows a moneyness region out to 0.94. Figure 9 extends this to 0.85. Again, there is little variation in the slope.

We continue to assume that all Brownian motions are mutually independent. The process for λ_t takes the same form as before, but instead of reverting to a constant value $\bar{\lambda}$, λ_t reverts to a value that is itself stochastic. We will assume that this value, ξ_t itself follows a square root process. Though the relative values of κ_λ and κ_ξ do not matter for the form of the solution, to enable an easier interpretation of these state variables, we will choose parameters so that shocks to ξ_t die out more slowly than direct shocks to λ_t. Duffie, Pan, and Singleton (2000) use the process in (13) and (14) to model return volatility. This model is also related to multi-factor return volatility processes proposed by Bates (2000) and Gallant, Hsu, and Tauchen (1999). In our study, the process is for jump intensity rather than volatility; however, volatility, which arises endogenously, will inherit the two-factor structure.

Like the one-factor SDR model, our two-factor extension is highly tractable. In Appendix C.1, we show that utility is given by

$$V_t = \frac{C_t^{1-\gamma}}{1-\gamma} e^{a + b_\lambda \lambda_t + b_\xi \xi_t} \tag{15}$$

where

$$a = \frac{1-\gamma}{\beta}\left(\mu - \frac{1}{2}\gamma\sigma^2\right) + \frac{b_\xi \kappa_\xi \bar{\xi}}{\beta} \tag{16}$$

$$b_\lambda = \frac{\kappa_\lambda + \beta}{\sigma_\lambda^2} - \sqrt{\left(\frac{\kappa_\lambda + \beta}{\sigma_\lambda^2}\right)^2 - 2\frac{E_\nu\left[e^{(1-\gamma)Z_t} - 1\right]}{\sigma_\lambda^2}} \tag{17}$$

$$b_\xi = \frac{\kappa_\xi + \beta}{\sigma_\xi^2} - \sqrt{\left(\frac{\kappa_\xi + \beta}{\sigma_\xi^2}\right)^2 - 2\frac{b_\lambda \kappa_\lambda}{\sigma_\xi^2}} \tag{18}$$

Equation 5 still holds for the state-price density, though the process V_t will differ. The riskfree rate and government bill rate are the same functions of λ_t as in the one-factor model. In Appendix C.3, we show that the price-dividend ratio is given by

$$G(\lambda_t, \xi_t) = \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda_t + b_{\phi\xi}(\tau)\xi_t\right), \tag{19}$$

where a_ϕ, $b_{\phi\lambda}$ and $b_{\phi\xi}$ solve the differential equations

$$\begin{aligned}
a'_\phi(\tau) &= -\beta - \mu - \gamma(\phi-1)\sigma^2 + \mu_D + b_{\phi\xi}(\tau)\kappa_\xi\bar{\xi} \\
b'_{\phi\lambda}(\tau) &= -b_{\phi\lambda}(\tau)\kappa_\lambda + \frac{1}{2}b_{\phi\lambda}(\tau)^2\sigma_\lambda^2 + b_\lambda b_{\phi\lambda}(\tau)\sigma_\lambda^2 + E_\nu\left[e^{(\phi-\gamma)Z_t} - e^{(1-\gamma)Z_t}\right] \\
b'_{\phi\xi}(\tau) &= -b_{\phi\lambda}(\tau)\kappa_\lambda - b_{\phi\xi}(\tau)\kappa_\xi + \frac{1}{2}b_{\phi\xi}(\tau)^2\sigma_\xi^2 + b_\xi b_{\phi\xi}(\tau)\sigma_\xi^2,
\end{aligned}$$

with boundary condition

$$a_\phi(0) = b_{\phi\lambda}(0) = b_{\phi\xi}(0) = 0.$$

Similar reasoning to that of Section 2.3 shows that for a given moneyness and time to expiration, normalized option prices and implied volatilities are a function of λ_t and ξ_t alone. To solve for option prices, we approximate $G(\lambda_t, \xi_t)$ by a log-linear function of λ_t and ξ_t, as shown in Appendix C.3.

4.3 Model results

4.3.1 Fitting the model

For simplicity, we keep risk aversion γ, the discount rate β and the leverage parameter ϕ the same as in the one-factor model. We also keep the distribution of consumption in the event of a disaster the same. Note that κ_λ and σ_λ will not have the same interpretation in the two-factor model as κ and σ_λ do in the one-factor model.

Our first goal in calibrating this new model is to generate reasonable predictions for the aggregate market and for the consumption distribution. That is, we do not want to allow the probability of a disaster to become too high. One challenge in calibrating representative agent models is to match the high volatility of the price-dividend ratio. In the two-factor model, as in the one-factor model, there is an upper limit to the amount of volatility that can be assumed in the state variable before a solution for utility fails to exist. The more persistent the processes, namely the lower the values of κ_λ and κ_ξ, the lower the respective

volatilities must be so as to ensure that the square root terms in (17) and (18) stay positive. We choose parameters keeping these various tradeoffs in mind, along with the restriction that $\kappa_\xi < \kappa_\lambda$.

The resulting parameter choices are shown in Table 2. The mean reversion parameter κ_λ and volatility parameter σ_λ are relatively high, indicating a fast-moving component to the λ_t process, while the mean reversion parameter κ_ξ and σ_ξ are relatively low, indicating a slower-moving component. The parameter $\bar{\xi}$ (which represents both the average value of ξ_t and the average value of λ_t) is 2% per annum. This is lower than $\bar{\lambda}$ in our calibration of the one-factor model. In this sense, the two-factor calibration is more conservative. However, the extra persistence created by the ξ_t process implies that λ_t could deviate from its average for long periods of time. To clarify the implications of these parameter choices, we report population statistics on λ_t in Panel C of Table 2. The median disaster probability is only 0.37%, indicating a highly skewed distribution. The standard deviation is 3.9% and the monthly first-order autocorrelation is 0.9858.

Implications for the riskfree rate and the market are shown in Table 3. We simulate 100,000 samples of length 60 years to capture features of the small-sample distribution. We also simulate a long sample of 600,000 years to capture the population distribution. Statistics are reported for the full set of 100,000 samples, and the subset for which there are no disasters (38% of the sample paths). The table reveals a good fit to the equity premium and to return volatility. The average Treasury Bill rate is slightly too high, though this could be lowered by lowering β or by lowering the probability of government default.[19] The model successfully captures the low volatility of the riskfree rate in the postwar period. The median value of the price-dividend ratio volatility is lower than in the data (0.27 versus 0.43), but the data value is still lower than the 95th percentile in the simulated sample.[20] For the market moments,

[19] As in the one-factor model, we assume a 40% probability of government default.

[20] The one-factor model, which was calibrated to match the population persistence of the price-dividend

only the very high AR(1) coefficient in postwar data falls outside the 90% confidence bounds: it is 0.92 (annual), while in the data, the median is 0.79 and the 95th percentile value is 0.91. As we will show below, there is a tension between matching the autocorrelation in the price-dividend ratio and in option prices. Moreover, so that utility converges, there is a tradeoff between persistence and volatility. One view is that the autocorrelation of the price-dividend ratio observed in the postwar period may in fact have been very exceptional and perhaps is not a moment that should be targeted too stringently.

4.3.2 Implications for the time series of option prices

Based on the discussion above, it seems likely that this model will, also, deliver a reasonable fit to the average three-month implied volatility curve. However, we ask more of this richer model; we also ask it to account for the time series of option prices.

We consider the time series of one-month ATM and OTM implied volatilities (Figure 8). For each of these data points, we compute the implied value of λ_t and ξ_t. Note that this exercise would not be possible if the model were not capable of simultaneously matching the level and slope of the implied volatility curve for one-month options. We show the resulting values in Figure 10. For most of the sample period, the disaster probability λ_t varies between 0 and 6%, with spikes corresponding to the Asian financial crisis in the late 1990s and the large market declines in the early 2000s. However, the sample is clearly dominated by the events of 2008-2009, in which the disaster probability rises to about 20%. While we choose the state variables to match the behavior of one-month options exactly, Figure 11 shows that the model also delivers a good fit to the time series of implied volatilities from three-month and six-month options.

The exercise above raises the question of whether the time series of λ_t and ξ_t are rea-

ratio, has a median price-dividend ratio volatility of 0.21 for sample paths without disasters and a population price-dividend ratio volatility of 0.38.

sonable given our assumptions on the processes for these variables. Table 4 calculates the distribution of moments for λ_t and ξ_t. With the exception of the first-order autoregressive coefficients, the data fall well within the 90% confidence intervals. That is, the average values of the state variables and their volatilities could easily have been observed in 17-year samples with no disasters. The persistence in the data is somewhat lower than the persistence implied by the model. This illustrates a tension between matching the time series of option prices and the time series of the price-dividend ratio, because while the model somewhat overstates the persistence in options, it understates the persistence in the price-dividend ratio. However, it is well known that the implied volatility series exhibits long-memory-like properties. Thus the first-order autocorrelation of volatility, and the implied λ_t and ξ_t may understate the true persistence of these variables in the data.

As discussed in Section 4.1, it is not clear why adding a second state variable would allow the model to capture the time series variation in both the level and the slope, since both of these observables would be endogenously determined by both state variables. To better understand the mechanisms that allow the model to match these data, we consider the relative contributions of two state variables to the implied volatility curve in Figure 12. Panel A of Figure 12 fixes ξ_t at its median value and shows implied volatilities for λ_t at its 80th percentile value, at its median, and at its 20th percentile value. Panel B of Figure 12 performs the analogous exercise, this time fixing λ_t but varying ξ_t. In contrast to its behavior in the one-factor model, increasing λ_t both increases implied volatilities and increases the slope. The state variable ξ_t has a smaller effect on the level of implied volatilities, but a greater effect on the slope. Moreover, an increase in ξ_t lowers the slope rather than raising it. To summarize, increases in λ_t raise both the level and the slope of the implied volatility curve. Increases in ξ_t slightly raise the level and decrease the slope.

Why does ξ_t affect the slope of the implied volatility curve? The reason is that ξ_t has a

relatively small affect on the probability of disasters at the time horizon important for option pricing, but a large effect on the volatility of stock prices (because of the square root term on its own volatility). Thus increases in ξ_t flatten the slope of the implied volatility curve because they raise the implied volatility for ATM options much more than for OTM options. The process for λ_t, on the other hand, has a large effect on OTM implied volatilities because it directly controls the probability of a disaster. It has a smaller effect on ATM implied volatilities than in the one-factor model because of its lower persistence.

4.3.3 The implied volatility surface

We now take a closer look at what the model says about the implied volatility surface, namely implied volatilities across moneyness and time to expiration. Figures 13–15 show the implied volatilities for one, three and six-month options as a function of moneyness. We show the average implied volatilities computed using the time series of λ_t and ξ_t, and compare these to the data. Given that we have chosen λ_t and ξ_t to match the time series of implied volatilities on 1-month options, it is not surprising that the model matches these data points exactly. More interestingly, the model is able to match the three-month and the six-month implied volatility curves almost exactly, even though it was not calibrated to these curves.

The downward slope in implied volatilities from options expiring in as long as six months indicates that the risk-neutral distribution of returns exhibits considerable skewness at long horizons. This is known as the skewness puzzle (Bates (2008)) because the law of large numbers would suggest convergence toward normality as the time to expiration increases. Recently, Neuberger (2012) makes use of options data to conclude that the skewness in the physical distribution of returns is also more pronounced than has been estimated previously. Neuberger emphasizes the observed negative correlation between stock prices and volatility (French, Schwert, and Stambaugh (1987), Schwert (1989)) as a reason why skewness in long-

horizon returns does not decay as the law of large numbers in an iid model suggests that it would (see also Bates (2000)).

Figure 15 shows that our model can capture the downward slope in 6-month implied volatilities as well as the slope for shorter-term options. Thus stock returns in the model exhibit skewness at both long and short-horizons. The short-horizon skewness arises from the existence of rare disasters. Long-horizon skewness, however, comes about endogenously because of the time-variation in the disaster probability. An increase in the rare disaster probability leads to lower stock prices, and, at the same time, higher volatilities, thereby accounting for this co-movement in the data. As a result, returns maintain their skewness at long horizons, and the model can explain six-month as well as one-month implied volatility curves.

5 Conclusion

Since the early work of Rubinstein (1994), the implied volatility curve has constituted an important piece of evidence against the Black-Scholes Model, and a lens through which to view the success of a model in matching option prices.

The implied volatility curve, almost by definition, has been associated with excess kurtosis in stock prices. Separately, a literature has developed linking kurtosis in consumption (which would then be inherited by returns in equilibrium) with the equity premium. However, much of the work up to now, as exemplified by a recent paper by Backus, Chernov, and Martin (2011) suggests that, at least for standard preferences, the non-normalities required to match the equity premium are qualitatively different from those required to match implied volatility.

We have proposed an alternative and more general approach to modeling the risk of downward jumps that can reconcile the implied volatility curve and the equity premium. Rather than assuming that the probability of a large negative event is constant, we allow it

to vary over time. The existence of very bad consumption events leads to both the downward slope in the implied volatility curve and the equity premium. However, the time-variation in these events moderates the slope, raises the level and generates the excess volatility observed in stock prices.

The initial model that we develop in the paper is deliberately simple and parsimonious. However, there are some interesting features of option and stock prices that cannot be matched by a model with a single state variable; for example, the imperfect correlation between the slope and the level of the implied volatility curve. For this reason, we investigate a more general model that allows for variation in disaster risk to occur at multiple time scales. This modification naturally produces time-variation in the slope of implied volatilities because it introduces variation in stock price volatility that can be distinguished from the risk of rare disasters. Taken together, these results indicate that options data support the existence of rare disasters in beliefs about the equity premium. Moreover, options data can provide information about the disaster distribution beyond that offered by stock prices. In particular, data from options suggest that modeling time-variation in disaster risk occurring at multiple time scales may be a fruitful avenue for future work.

Appendix

A The state-price density

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

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

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

Equation (A.1) shows the state-price density can be expressed in terms of a locally deterministic term and a term that is locally stochastic. To obtain (12), we require both to be expressed in terms of C_t and λ_t. We derive the expression for the stochastic term first.

It follows from (4) that

$$\frac{\partial}{\partial C} f\left(C_t, V_t\right) = \beta\left(1-\gamma\right)\frac{V_t}{C_t}. \tag{A.2}$$

Wachter (2013) shows that continuation utility V_t can be expressed in terms of C_t as follows:[21]

$$V_t = J(C_t, \lambda_t), \tag{A.3}$$

where

$$J(C_t, \lambda_t) = \frac{C_t^{1-\gamma}}{1-\gamma} e^{a+b\lambda_t}, \tag{A.4}$$

and

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

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

[21] Wachter (2013) expresses the value function in terms of wealth rather than consumption. Because the ratio of wealth to consumption is β^{-1}, it is straightforward to go from one expression to the other. The difference between the expressions is found in the definition of a. In the earlier paper, this expression has an extra term given by $(1-\gamma)\log\beta$.

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

$$\frac{1}{2}\sigma_\lambda^2 b^2 - (\kappa+\beta)b + E_\nu\left[e^{(1-\gamma)Z} - 1\right] = 0. \tag{A.7}$$

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

$$\frac{\partial}{\partial C} f\left(C_t, V_t\right) = \beta C_t^{-\gamma} e^{a+b\lambda_t}. \tag{A.8}$$

It also follows from (4) that

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

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

$$\frac{\partial}{\partial V} f(C_t, V_t) = -\beta(a + b\lambda_t) - \beta. \tag{A.9}$$

Finally, we collect constant terms:

$$\eta = -\beta a - \beta \tag{A.10}$$

so that

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

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

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

A.2 The iid limit of the SDR model.

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

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

L'Hopital's rule implies

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

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

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

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

We now apply these results to calculate the limit of π_t/π_0. It follows from (A.1), (A.8) and (A.9) that

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

which is equivalent to the result one obtains by calculating the state price density in the iid case. Note that this result is not automatic, but rather holds because we choose the lower of the two roots of (A.7).[22]

A.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.

[22] This point is also made by Tauchen (2005) for a model with stochastic volatility.

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 A.2,

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

It follows that, for β given by

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

ratios of the state price densities are the same.

B Details of the calculation of option prices

B.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

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

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_t close to λ^*,

$$g(\lambda) \simeq g(\lambda^*) + (\lambda - \lambda^*)g'(\lambda^*). \tag{B.1}$$

Moreover,

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

The expression (B.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

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

and the log-linear function

$$\hat{G}(\lambda) = G(\lambda^*) \exp\left\{b_\phi^*(\lambda - \lambda^*)\right\}. \tag{B.4}$$

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

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

When we apply this method in this paper, we choose λ^* equal to the long-run mean $\bar{\lambda}$.

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.

Figure B.1 shows implied volatilities from option prices computed using the loglinear approximation described above, and from option prices computed by solving the expectation in (10) directly, using by averaging over simulated sample paths. To keep the computation tractable, we assume a single jump size of -30%. The implied volatilities are extremely close in the two cases.

B.2 Transform analysis

The normalized put option price is given as

$$P^n(\lambda_t, T-t; K^n) = E_t\left[\frac{\pi_T}{\pi_t}\left(K^n - \frac{F_T}{F_t}\right)^+\right]. \tag{B.5}$$

It follows from (11), (12), and (B.4) that

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

where η, b and b^*_ϕ are constants defined by (A.10), (7) and (B.3), respectively. Then (B.5) can be rewritten as

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

Note that

$$1_{\left\{\frac{F_T}{F_t}\le K^n\right\}} = 1_{\left\{b^*_\phi(\lambda_T-\lambda_t)+\phi(\log C_T-\log C_t)\le\log K^n\right\}}.$$

Equation (B.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.

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

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

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

and let

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

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

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

and therefore

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

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.

C Solving the two-factor stochastic disaster risk model

C.1 Utility

We conjecture that the value function is given by

$$J(C, \lambda, \xi) = \frac{C^{1-\gamma}}{1-\gamma} e^{a + b_\lambda \lambda + b_\xi \xi}. \tag{C.1}$$

It follows from the form of $f(C, V)$ that

$$\begin{aligned} f(C, V) &= \beta(1-\gamma)V\left(\log C - \frac{1}{1-\gamma}\log\left[(1-\gamma)V\right]\right) \\ &= \beta(1-\gamma)V\log C - \beta V \log\left[(1-\gamma)V\right] \\ &= \beta V \log\left(\frac{C^{1-\gamma}}{(1-\gamma)V}\right) \\ &= -\beta V(a + b_\lambda \lambda + b_\xi \xi), \end{aligned}$$

where the last equation follows from the equilibrium condition that the utility process is equal to the value function under the optimal policies: $V_t = J(C_t, \lambda_t, \xi_t)$.

By differentiating $J(C, \lambda, \xi)$, we obtain

$$\begin{aligned}
\frac{\partial J}{\partial C} &= (1-\gamma)\frac{J}{C}, & \frac{\partial^2 J}{\partial C^2} &= -\gamma(1-\gamma)\frac{J}{C^2}, \\
\frac{\partial J}{\partial \lambda} &= b_\lambda J, & \frac{\partial^2 J}{\partial \lambda^2} &= b_\lambda^2 J, \\
\frac{\partial J}{\partial \xi} &= b_\xi J, & \frac{\partial^2 J}{\partial \xi^2} &= b_\xi^2 J.
\end{aligned} \tag{C.2}$$

Applying Ito's Lemma to $J(C, \lambda, \xi)$ with conjecture (C.1) and derivatives (C.2):

$$\begin{aligned}
\frac{dV_t}{V_{t^-}} = (1-\gamma)(\mu dt + \sigma dB_t) - \frac{1}{2}\gamma(1-\gamma)\sigma^2 dt \\
+ b_\lambda \left(\kappa_\lambda(\xi_t - \lambda_t)dt + \sigma_\lambda \sqrt{\lambda_t} dB_{\lambda,t}\right) + \frac{1}{2} b_\lambda^2 \sigma_\lambda^2 \lambda_t dt \\
+ b_\xi \left(\kappa_\xi(\bar{\xi} - \xi_t)dt + \sigma_\xi \sqrt{\xi_t} dB_{\xi,t}\right) + \frac{1}{2} b_\xi^2 \sigma_\xi^2 \xi_t dt + (e^{(1-\gamma)Z_t} - 1)dN_t.
\end{aligned}$$

Under the optimal consumption path, it must be that

$$V_t + \int_0^t f(C_s, V_s)ds = E_t\left[\int_0^\infty f(C_s, V_s)ds\right] \tag{C.3}$$

(see Duffie and Epstein (1992)). By the law of iterative expectations, the left-hand side of (C.3) is a martingale. Thus, the sum of the drift and the jump compensator of $(V_t + \int_0^t f(C_s, V_s)ds)$ equals zero. That is,

$$\begin{aligned}
0 = (1-\gamma)\mu - \frac{1}{2}\gamma(1-\gamma)\sigma^2 + b_\lambda \kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2} b_\lambda^2 \sigma_\lambda^2 \lambda_t + b_\xi \kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2} b_\xi^2 \sigma_\xi^2 \xi_t \\
+ \lambda_t E_\nu\left[e^{(1-\gamma)Z_t} - 1\right] - \beta(a + b_\lambda \lambda_t + b_\xi \xi_t).
\end{aligned} \tag{C.4}$$

By collecting terms in (C.4), we obtain

$$0 = \underbrace{\left[(1-\gamma)\mu - \frac{1}{2}\gamma(1-\gamma)\sigma^2 + b_\xi \kappa_\xi \bar{\xi} - \beta a\right]}_{=0}$$
$$+ \lambda_t \underbrace{\left[-b_\lambda \kappa_\lambda + \frac{1}{2} b_\lambda^2 \sigma_\lambda^2 + E_\nu\left[e^{(1-\gamma)Z_t} - 1\right] - \beta b_\lambda\right]}_{=0}$$
$$+ \xi_t \underbrace{\left[b_\lambda \kappa_\lambda - b_\xi \kappa_\xi + \frac{1}{2} b_\xi^2 \sigma_\xi^2 - \beta b_\xi\right]}_{=0}.$$

Solving these equations gives us

$$\begin{aligned}
a &= \frac{1-\gamma}{\beta}\left(\mu - \frac{1}{2}\gamma\sigma^2\right) + \frac{b_\xi \kappa_\xi \bar{\xi}}{\beta} \\
b_\lambda &= \frac{\kappa_\lambda + \beta}{\sigma_\lambda^2} - \sqrt{\left(\frac{\kappa_\lambda + \beta}{\sigma_\lambda^2}\right)^2 - 2\frac{E_\nu\left[e^{(1-\gamma)Z_t} - 1\right]}{\sigma_\lambda^2}} \\
b_\xi &= \frac{\kappa_\xi + \beta}{\sigma_\xi^2} - \sqrt{\left(\frac{\kappa_\xi + \beta}{\sigma_\xi^2}\right)^2 - 2\frac{b_\lambda \kappa_\lambda}{\sigma_\xi^2}},
\end{aligned}$$

where we have chosen the negative root based on the economic consideration that when there are no disasters, λ_t and ξ_t should not appear in the value function. Namely, for $Z_t = 0$, $b_\lambda = b_\xi = 0$. Lastly, note that these results verify the conjecture (C.1).

C.2 State-price density

Following the same steps for the SDR model, we can show that

$$\begin{aligned}
\frac{\partial}{\partial C} f(C_t, V_t) &= \beta C_t^{-\gamma} e^{a + b_\lambda \lambda_t + b_\xi \xi_t} \\
\frac{\partial}{\partial V} f(C_t, V_t) &= \eta - \beta b_\lambda \lambda_t - \beta b_\xi \xi_t.
\end{aligned}$$

Therefore, it follows that the state-price density can be written as

$$\pi_t = \exp\left(\eta t - \beta b_\lambda \int_0^t \lambda_s ds - \beta b_\xi \int_0^t \xi_s ds\right) \beta C_t^{-\gamma} e^{a + b_\lambda \lambda_t + b_\xi \xi_t}. \tag{C.5}$$

By applying Ito's Lemma to (C.5), we derive the following stochastic differential equation:

$$\frac{d\pi_t}{\pi_{t^-}} = \left\{-\beta - \mu + \gamma\sigma^2 - \lambda_t E_\nu \left[e^{(1-\gamma)Z_t} - 1\right]\right\} dt$$
$$- \gamma\sigma dB_t + b_\lambda \sigma_\lambda \sqrt{\lambda_t} dB_{\lambda,t} + b_\xi \sigma_\xi \sqrt{\xi_t} dB_{\xi,t} + (e^{-\gamma Z_t} - 1) dN_t.$$

At equilibrium, the sum of the drift and the jump compensator of the state-price density growth must equal the riskfree rate multiplied by -1. It thus follows that the riskfree rate is given by

$$r_{f,t} = \beta + \mu - \gamma\sigma^2 + \lambda_t E_\nu \left[e^{(1-\gamma)Z_t} - e^{-\gamma Z_t}\right].$$

C.3 Dividend claim price

Let F_t denote the price of the dividend claim. The pricing relation implies

$$\begin{aligned} F_t &= E_t \left[\int_t^\infty \frac{\pi_s}{\pi_t} D_s ds\right] \\ &= \int_t^\infty E_t \left[\frac{\pi_s}{\pi_t} D_s\right] ds. \end{aligned}$$

Let $H(D_t, \lambda_t, \xi_t, s-t)$ denote the price of the asset that pays the aggregate dividend at time s, namely,

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

By the law of iterative expectations, it follows that $\pi_t H_t$ is a martingale:

$$\pi_t H(D_t, \lambda_t, \xi_t, s-t) = E_t[\pi_s D_s].$$

Conjecture that

$$H(D_t, \lambda_t, \xi_t, \tau) = D_t \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda_t + b_{\phi\xi}(\tau)\xi_t\right). \tag{C.6}$$

Ito's Lemma implies

$$\frac{dH_t}{H_{t-}} = \Big\{\mu_D + b_{\phi\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{\phi\lambda}(\tau)^2\sigma_\lambda^2\lambda_t + b_{\phi\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{\phi\xi}(\tau)^2\sigma_\xi^2\xi_t$$
$$- a'_\phi(\tau) - b'_{\phi\lambda}(\tau)\lambda_t - b'_{\phi\xi}(\tau)\xi_t\Big\}\,dt$$
$$+ \phi\sigma dB_t + b_{\phi\lambda}(\tau)\sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} + b_{\phi\xi}(\tau)\sigma_\xi\sqrt{\xi_t}dB_{\xi,t} + (e^{\phi Z_t} - 1)dN_t.$$

Combining the SDE for H_t with the one for π_t derived in the previous sections, we can derive the SDE for $\pi_t H_t$:

$$\frac{d(\pi_t H_t)}{\pi_{t^-}H_{t^-}} = \Big\{ - \beta - \mu + \gamma\sigma^2 - \lambda_t E_\nu\left[e^{(1-\gamma)Z_t} - 1\right]$$
$$+ \mu_D + b_{\phi\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{\phi\lambda}(\tau)^2\sigma_\lambda^2\lambda_t$$
$$+ b_{\phi\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{\phi\xi}(\tau)^2\sigma_\xi^2\xi_t$$
$$- a'_\phi(\tau) - b'_{\phi\lambda}(\tau)\lambda_t - b'_{\phi\xi}(\tau)\xi_t$$
$$- \gamma\phi\sigma^2 + b_\lambda b_{\phi\lambda}(\tau)\sigma_\lambda^2\lambda_t + b_\xi b_{\phi\xi}(\tau)\sigma_\xi^2\xi_t\Big\}\,dt$$
$$+ (\phi - \gamma)\sigma dB_t + (b_\lambda + b_{\phi\lambda}(\tau))\sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} + (b_\xi + b_{\phi\xi}(\tau))\sigma_\xi\sqrt{\xi_t}dB_{\xi,t}$$
$$+ (e^{(\phi-\gamma)Z_t} - 1)dN_t.$$

Since $\pi_t H_t$ is a martingale, the sum of the drift and the jump compensator of $\pi_t H_t$ equals zero. Thus:

$$0 = -\beta - \mu + \gamma\sigma^2 - \lambda_t E_\nu\left[e^{(1-\gamma)Z_t} - 1\right]$$
$$+ \mu_D + b_{\phi\lambda}(\tau)\kappa_\lambda(\xi_t - \lambda_t) + \frac{1}{2}b_{\phi\lambda}(\tau)^2\sigma_\lambda^2\lambda_t$$
$$+ b_{\phi\xi}(\tau)\kappa_\xi(\bar{\xi} - \xi_t) + \frac{1}{2}b_{\phi\xi}(\tau)^2\sigma_\xi^2\xi_t$$
$$- a'_\phi(\tau) - b'_{\phi\lambda}(\tau)\lambda_t - b'_{\phi\xi}(\tau)\xi_t$$
$$- \gamma\phi\sigma^2 + b_\lambda b_{\phi\lambda}(\tau)\sigma_\lambda^2\lambda_t + b_\xi b_{\phi\xi}(\tau)\sigma_\xi^2\xi_t + \lambda_t E_\nu\left[e^{(\phi-\gamma)Z_t} - 1\right]. \quad \text{(C.7)}$$

Collecting terms of (C.7) results in the following equation:

$$
\begin{aligned}
0 = & \underbrace{\left[-\beta-\mu+\gamma\sigma^2+\mu_D+b_{\phi\xi}(\tau)\kappa_\xi\bar{\xi}-\gamma\phi\sigma^2-a'_\phi(\tau)\right]}_{=0} \\
& + \lambda_t \underbrace{\left[-b_{\phi\lambda}(\tau)\kappa_\lambda+\frac{1}{2}b_{\phi\lambda}(\tau)^2\sigma_\lambda^2+b_\lambda b_{\phi\lambda}(\tau)\sigma_\lambda^2+E_\nu\left[e^{(\phi-\gamma)Z_t}-e^{(1-\gamma)Z_t}\right]-b'_{\phi\lambda}(\tau)\right]}_{=0} \\
& + \xi_t \underbrace{\left[b_{\phi\lambda}(\tau)\kappa_\lambda-b_{\phi\xi}(\tau)\kappa_\xi+\frac{1}{2}b_{\phi\xi}(\tau)^2\sigma_\xi^2+b_\xi b_{\phi\xi}(\tau)\sigma_\xi^2-b'_{\phi\xi}(\tau)\right]}_{=0}.
\end{aligned}
$$

It follows that

$$
\begin{aligned}
a'_\phi(\tau) &= -\beta-\mu-\gamma(\phi-1)\sigma^2+\mu_D+b_{\phi\xi}(\tau)\kappa_\xi\bar{\xi} \\
b'_{\phi\lambda}(\tau) &= -b_{\phi\lambda}(\tau)\kappa_\lambda+\frac{1}{2}b_{\phi\lambda}(\tau)^2\sigma_\lambda^2+b_\lambda b_{\phi\lambda}(\tau)\sigma_\lambda^2+E_\nu\left[e^{(\phi-\gamma)Z_t}-e^{(1-\gamma)Z_t}\right] \\
b'_{\phi\xi}(\tau) &= b_{\phi\lambda}(\tau)\kappa_\lambda-b_{\phi\xi}(\tau)\kappa_\xi+\frac{1}{2}b_{\phi\xi}(\tau)^2\sigma_\xi^2+b_\xi b_{\phi\xi}(\tau)\sigma_\xi^2. \qquad \text{(C.8)}
\end{aligned}
$$

This establishes that H satisfies the conjecture (C.6). We note that by no-arbitrage,

$$H(D_t,\lambda_t,\xi_t,0)=D_t$$

This condition provides the boundary conditions for the system of ODEs (C.8):

$$a_\phi(0)=b_{\phi\lambda}(0)=b_{\phi\xi}(0)=0$$

Recall that once we get $a_\phi(\tau)$, $b_{\phi\lambda}(\tau)$, and $b_{\phi\xi}(\tau)$,

$$
\begin{aligned}
F_t &= \int_t^\infty E_t\left[\frac{\pi_s}{\pi_t}D_s\right]ds \\
&= \int_t^\infty H(D_t,\lambda_t,\xi_t,s-t)ds \\
&= D_t\int_t^\infty \exp\left(a_\phi(s-t)+b_{\phi\lambda}(s-t)\lambda_t+b_{\phi\xi}(s-t)\xi_t\right)ds \\
&= D_t\int_0^\infty \exp\left(a_\phi(\tau)+b_{\phi\lambda}(\tau)\lambda_t+b_{\phi\xi}(\tau)\xi_t\right)d\tau.
\end{aligned}
$$

That is, the price-dividend ratio can be written as

$$G(\lambda_t, \xi_t) = \int_0^\infty \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda_t + b_{\phi\xi}(\tau)\xi_t\right) d\tau.$$

C.4 Approximating the price-dividend ratio

Let $g(\lambda, \xi) = \log G(\lambda, \xi)$. For given λ^* and ξ^*, the two-dimensional Taylor approximation implies

$$g(\lambda, \xi) \simeq g(\lambda^*, \xi^*) + \left.\frac{\partial g}{\partial \lambda}\right|_{\lambda^*, \xi^*} (\lambda - \lambda^*) + \left.\frac{\partial g}{\partial \xi}\right|_{\lambda^*, \xi^*} (\xi - \xi^*). \tag{C.9}$$

We note that

$$\begin{aligned}
\left.\frac{\partial g}{\partial \lambda}\right|_{\lambda^*, \xi^*} &= \frac{1}{G(\lambda^*, \xi^*)} \left.\frac{\partial G}{\partial \lambda}\right|_{\lambda^*, \xi^*} \\
&= \frac{1}{G(\lambda^*, \xi^*)} \int_0^\infty b_{\phi\lambda}(\tau) \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda^* + b_{\phi\xi}(\tau)\xi^*\right) d\tau.
\end{aligned}$$

Similarly, we obtain

$$\begin{aligned}
\left.\frac{\partial g}{\partial \xi}\right|_{\lambda^*, \xi^*} &= \frac{1}{G(\lambda^*, \xi^*)} \left.\frac{\partial G}{\partial \xi}\right|_{\lambda^*, \xi^*} \\
&= \frac{1}{G(\lambda^*, \xi^*)} \int_0^\infty b_{\phi\xi}(\tau) \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda^* + b_{\phi\xi}(\tau)\xi^*\right) d\tau.
\end{aligned}$$

We define the notation

$$\begin{aligned}
b^*_{\phi\lambda} &= \frac{1}{G(\lambda^*, \xi^*)} \int_0^\infty b_{\phi\lambda}(\tau) \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda^* + b_{\phi\xi}(\tau)\xi^*\right) d\tau \\
b^*_{\phi\xi} &= \frac{1}{G(\lambda^*, \xi^*)} \int_0^\infty b_{\phi\xi}(\tau) \exp\left(a_\phi(\tau) + b_{\phi\lambda}(\tau)\lambda^* + b_{\phi\xi}(\tau)\xi^*\right) d\tau,
\end{aligned}$$

and the log-linear function

$$\hat{G}(\lambda, \xi) = G(\lambda^*, \xi^*) \exp\left\{b^*_{\phi\lambda}(\lambda - \lambda^*) + b^*_{\phi\xi}(\xi - \xi^*)\right\}.$$

It follows from exponentiating both sides of (C.9) that

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

In our analysis, we pick λ^* and ξ^* as the stationary mean of λ_t and ξ_t, respectively.

References

At-Sahalia, Yacine, Yubo Wang, and Francis Yared, 2001, Do option markets correctly price the probabilities of movement of the underlying asset?, *Journal of Econometrics* 102, 67–110.

Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters Implied by Equity Index Options, *The Journal of Finance* 66, 1969–2012.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical Performance of Alternative Option Pricing Models, *The Journal of Finance* 52, 2003–2049.

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 Jose F. Ursua, 2008, Macroeconomic crises since 1870, *Brookings Papers on Economic Activity* no. 1, 255–350.

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.

Black, Fischer, and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637–654.

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected Stock Returns and Variance Risk Premia, *Review of Financial Studies* 22, 4463–4492.

Bollerslev, Tim, and Viktor Todorov, 2011, Tails, Fears, and Risk Premia, *The Journal of Finance* 66, 2165–2211.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premia: Evidence from futures options, *Journal of Finance* 62, 1453–1490.

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.

Carr, Peter, and Liuren Wu, 2003, What Type of Process Underlies Options? A Simple Robust Test, *The Journal of Finance* 58, 2581–2610.

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, Scott Joslin, and Ngoc-Khanh Tran, 2012, Rare Disasters and Risk Sharing with Heterogeneous Beliefs, *Review of Financial Studies* 25, 2189–2224.

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.

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, 2012, Uncertainty, Time-Varying Fear, and Asset Prices, forthcoming, *Journal of Finance*.

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, Bjorn, and Ivan Shaliastovich, 2008, An equilibrium guide to designing affine pricing models, *Mathematical Finance* 18, 519–543.

Fisher, Mark, and Christian Gilles, 1999, Consumption and asset prices with homothetic recursive preferences, Working paper, 99-17 Federal Reserve Bank of Atlanta.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

Gabaix, Xavier, 2008, Linearity-generating processes: A modelling tool yielding closed forms for asset prices, Working paper, New York University.

Gabaix, Xavier, 2012, An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.

Gallant, A. Ronald, Chien-Te Hsu, and George Tauchen, 1999, Using Daily Range Data to Calibrate Volatility Diffusions and Extract the Forward Integrated Variance, *Review of Economics and Statistics* 81, 617–631.

Gourio, Francois, 2011, Disaster Risk and Business Cycles, Forthcoming, *American Economic Review*.

Kocherlakota, Narayana R., 1996, The Equity Premium: It's Still a Puzzle, *Journal of Economic Literature* 34, 42–71.

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.

Mehra, Rajnish, and Edward Prescott, 1985, The equity premium puzzle, *Journal of Monetary Economics* 15, 145–161.

Mehra, Rajnish, and Edward C. Prescott, 2003, The equity premium in retrospect, in G. M. Constantinides, M. Harris, and R. M. Stulz, eds.: *Handbook of the Economics of Finance* (Elsevier, North-Holland).

Merton, Robert C., 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.

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.

Rietz, Thomas A., 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117–131.

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.

Schroder, Mark, and Costis Skiadas, 1999, Optimal consumption and portfolio selection with stochastic differential utility, *Journal of Economic Theory* 89, 68–126.

Schwert, William G., 1989, Why does stock market volatility change over time?, *Journal of Finance* 44, 1115–1153.

Shaliastovich, Ivan, 2009, Learning, Confidence and Option Prices, working paper, University of Pennsylvania.

Siegel, Jeremy J., 1994, *Stocks for the long run: a guide to selecting markets for long-term growth.* (Irwin Burr Ridge, IL).

Tauchen, George, 2005, Stochastic volatility in general equilibrium, Working paper, Duke University.

Veronesi, Pietro, 2004, The Peso problem hypothesis and stock market returns, *Journal of Economic Dynamics and Control* 28, 707–725.

Wachter, Jessica A., 2013, Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?, *The Journal of Finance* 68, 987–1035.

Weil, Philippe, 1990, Nonexpected utility in macroeconomics, *Quarterly Journal of Economics* 105, 29–42.

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

Yan, Shu, 2011, Jump risk, stock returns, and slope of implied volatility smile, *Journal of Financial Economics* 99, 216–233.

Figure 1: Histogram of consumption declines in the stochastic disaster risk model

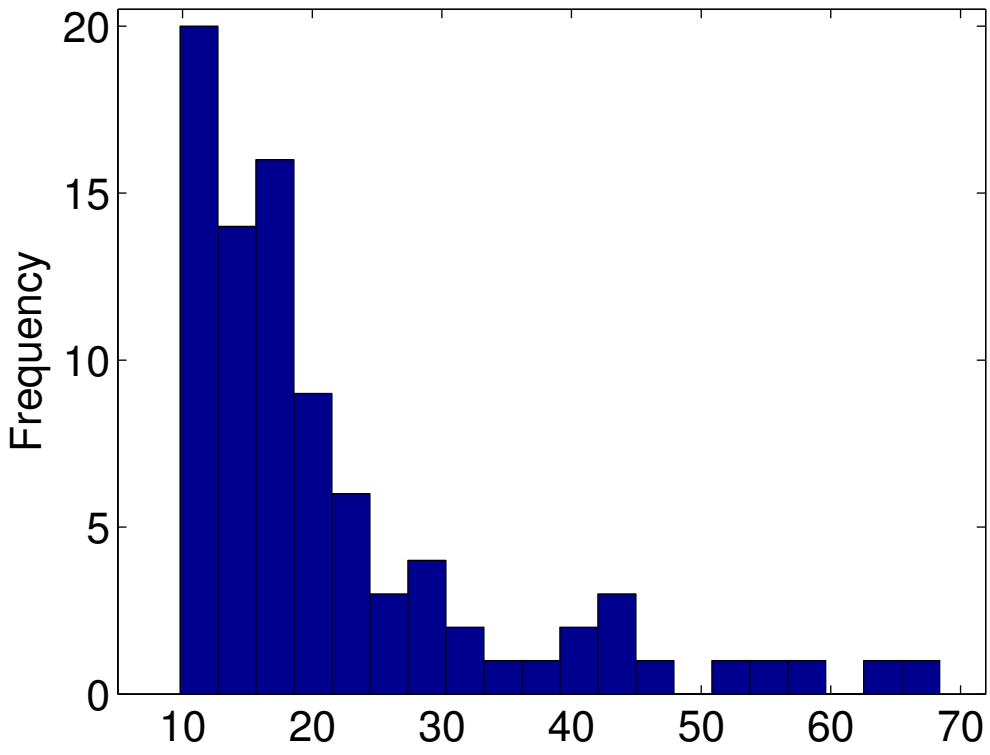

Notes: Histogram of large consumption declines (in percentages). The distribution is calculated by Barro and Ursua (2008) based on a century-long dataset of 22 countries. The histogram is for the quantity $1 - e^Z$ in the model.

Figure 2: Probability density functions for consumption declines

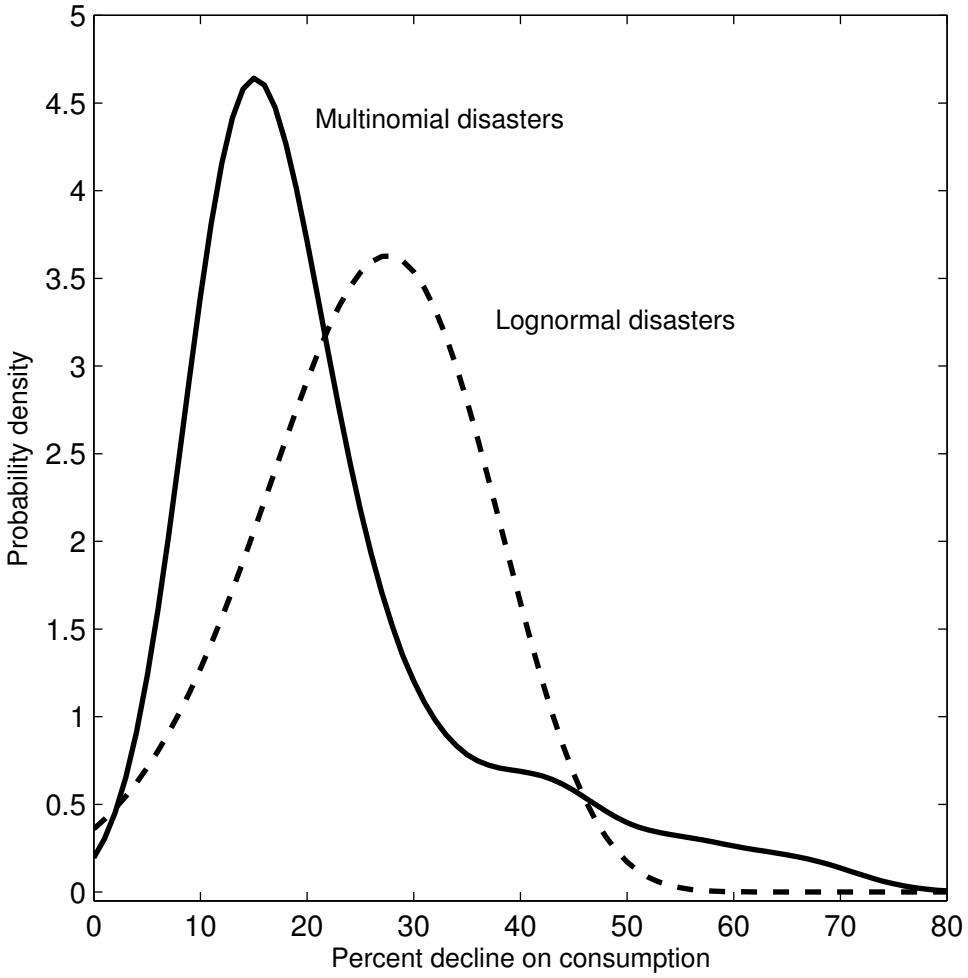

Notes: The probability density functions (pdfs) for consumption declines for log-normally distributed disasters and for the multinomial distribution assumed in the stochastic disaster risk (SDR) model. In the case of the SDR model, the pdf approximates the multinomial distribution from Barro and Ursua (2008). The exact multinomial distribution is used to calculate the results in the paper. The pdfs are for the quantities $1 - e^Z$ in each model.

Figure 3: Average implied volatilities in the SDR and CDR models

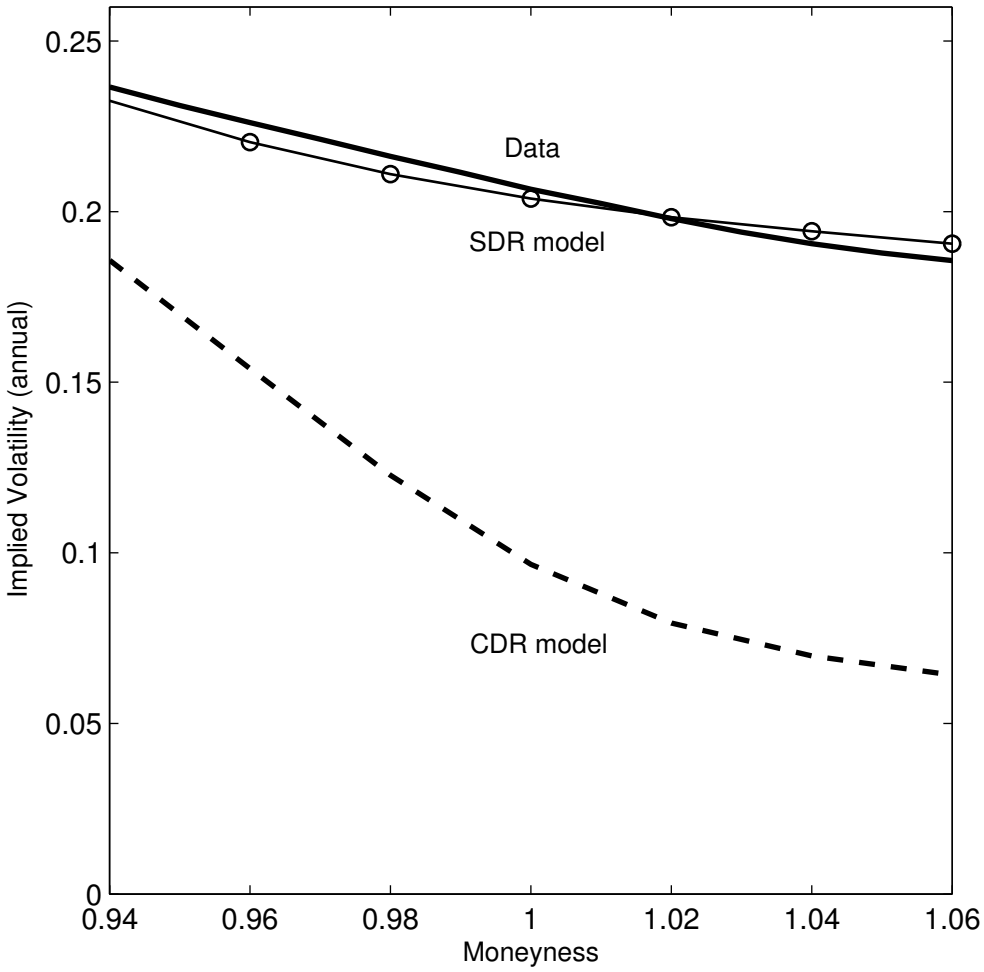

Notes: Average implied volatilities for 3-month options as a function of moneyness for the stochastic disaster risk (SDR) model, for the constant disaster risk (CDR) model (under the benchmark calibration given in Table 1) and in the data. Average implied volatilities are shown as functions of moneyness, defined as the exercise price divided by the asset price.

Figure 4: Implied volatilities as a function of the disaster probability for the SDR model

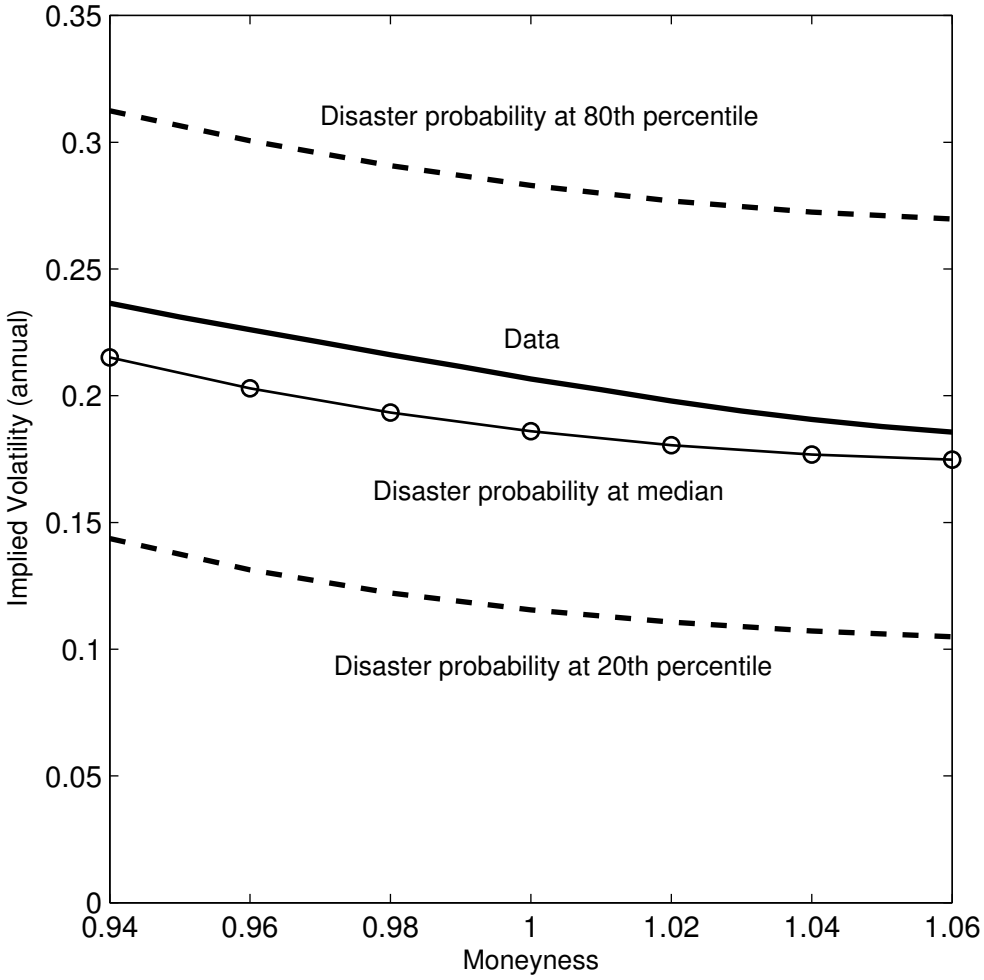

Notes: Implied volatilities on 3-month options as a function of moneyness and of the probability of a disaster for the SDR model.

Figure 5: Comparative statics for the CDR model

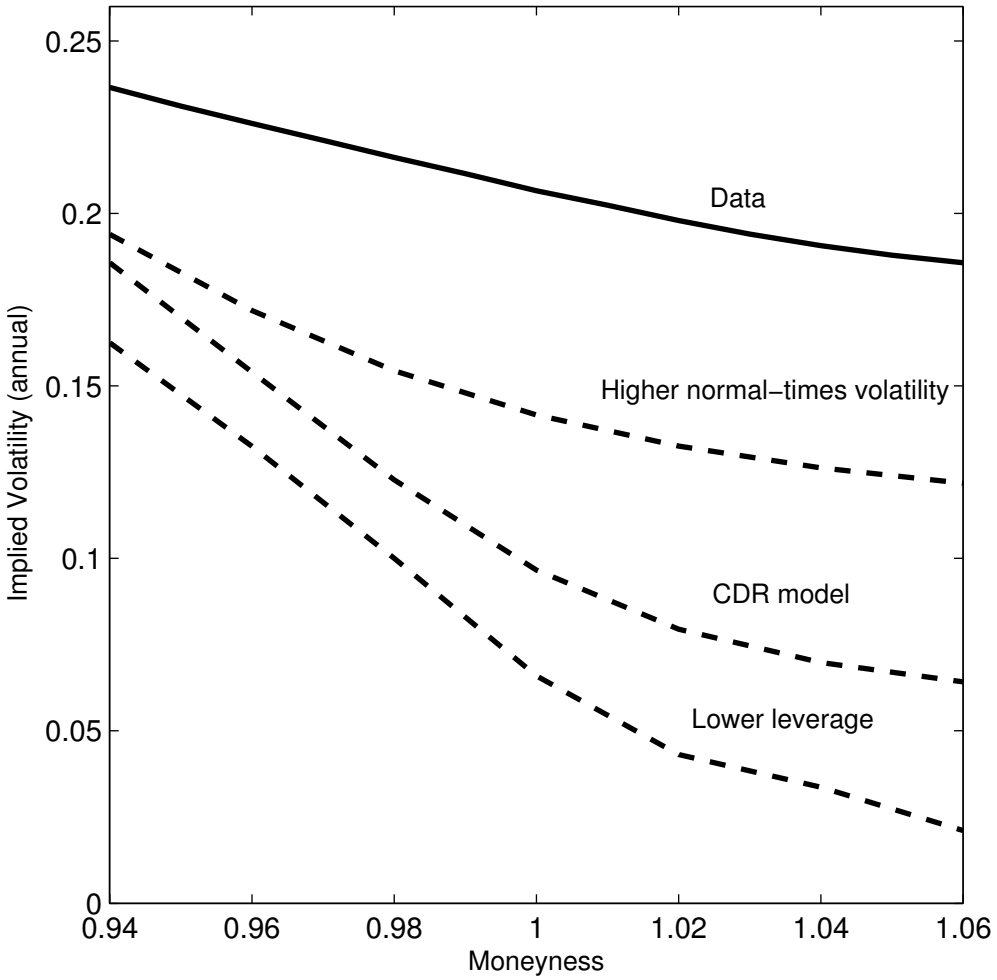

Notes: Implied volatilities for 3-month options as a function of moneyness in the data and for three parameterization of the CDR model. The line labeled "CDR" shows the benchmark calibration. The line labeled "higher normal-times volatility" raises the volatility of consumption shocks that are not associated with disasters from 1% to 2% per annum but keeps all other parameters, including the consumption disaster distribution, the same. The line labeled "lower leverage" lowers the term multiplying dividends from 5.1 to 2.6, while keeping all other parameters the same.

Figure 6: Average implied volatilities in the SDR model and in an alternative calibration of the CDR model

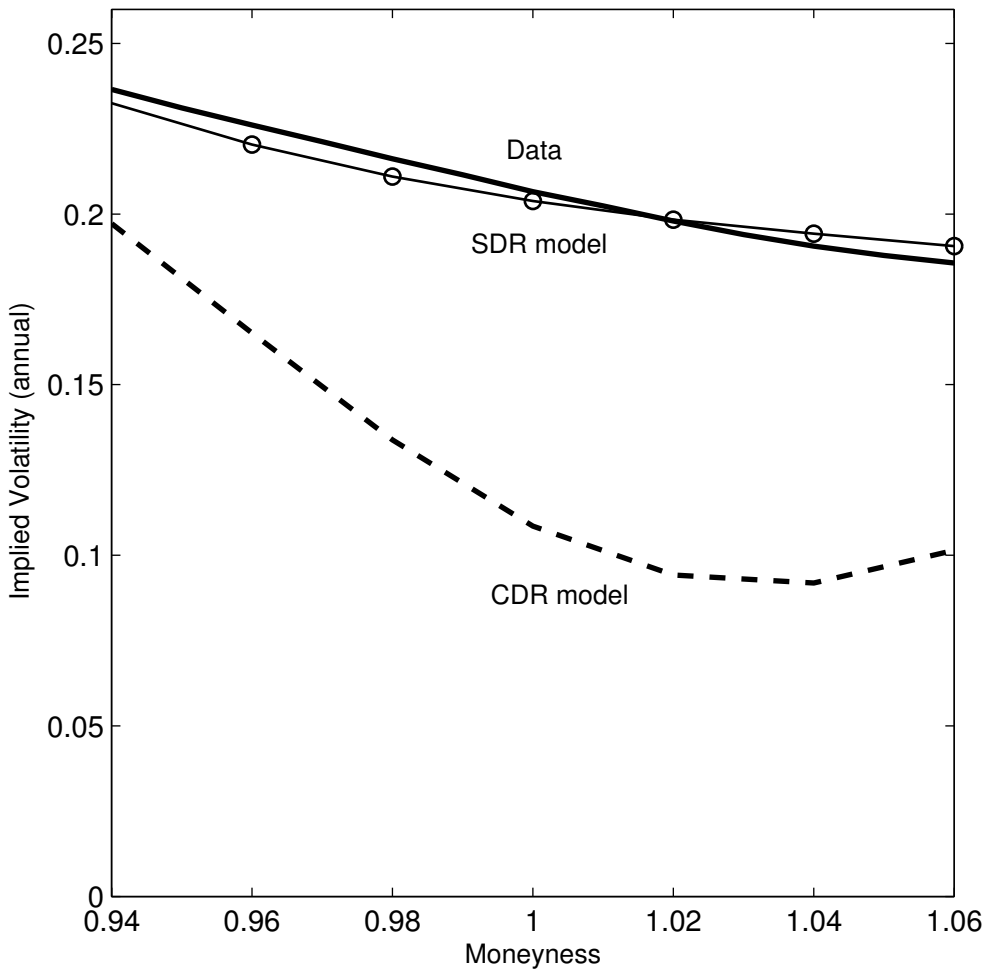

Notes: Implied volatilities for 3-month options as a function of moneyness for the stochastic disaster risk (SDR) model, for the constant disaster risk (CDR) model and in the data. The parameters for the CDR model are chosen to be identical to those of the SDR model, except that σ_λ and κ are set to zero.

Figure 7: Evaluating the role of recursive utility

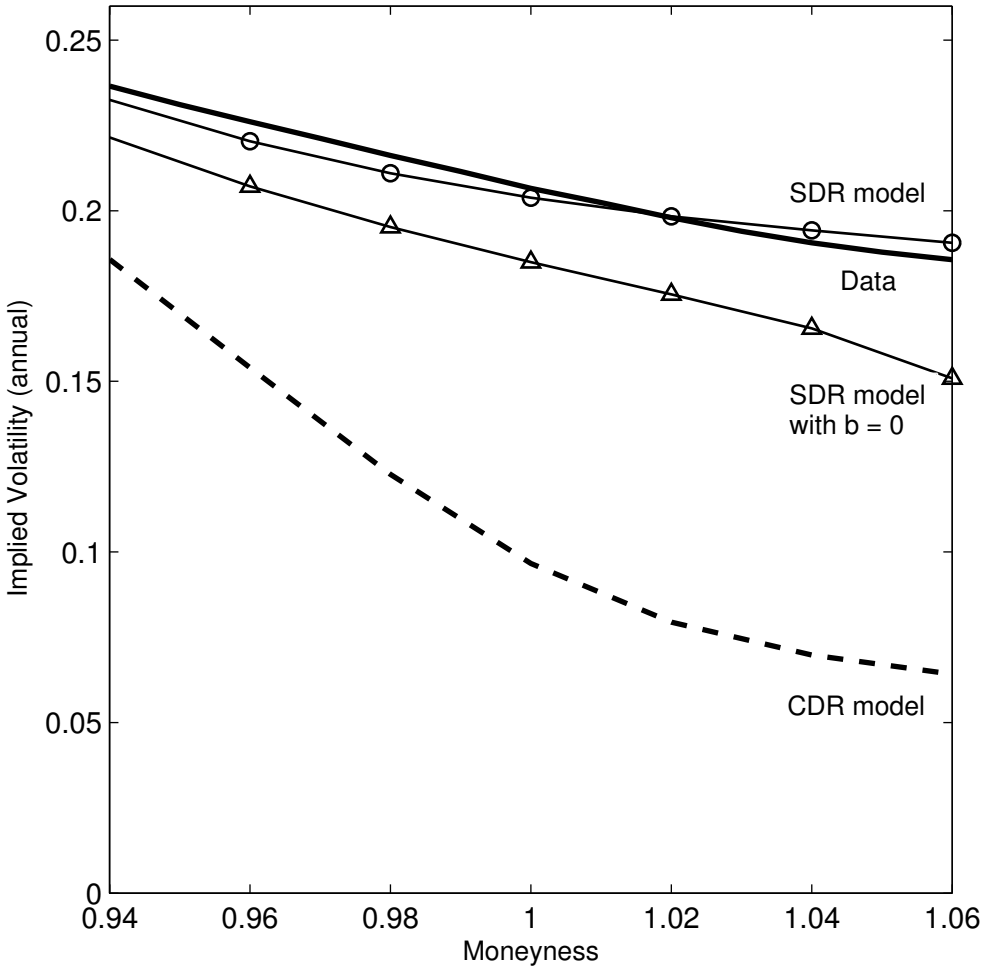

Notes: Implied volatilities for 3-month options as a function of moneyness in the data, in the CDR model (calibrated as in Table 1) and in the SDR model. Also shown are implied volatilities in the SDR model computed under the assumption that the premium associated with time-variation in the disaster probability is equal to zero (SDR model with $b = 0$).

Figure 8: 1-month implied volatility time series

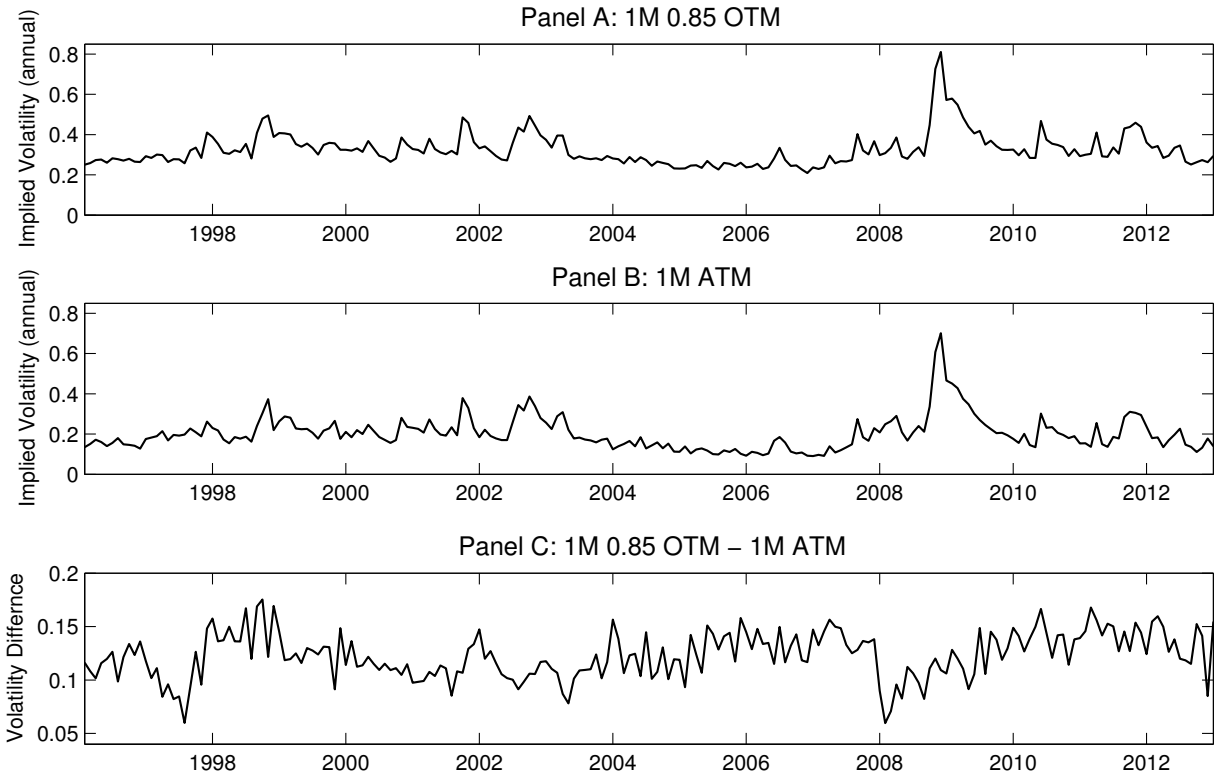

Notes: Monthly time series of implied volatilities on 1-month options in the data and of the difference in implied volatilities. Implied volatilities are calculated for at-the-money (ATM) options and out-of-the-money (OTM) options with moneyness equal to 0.85.

Figure 9: Implied volatilities as a function of the disaster probability in the SDR model for extended moneyness range

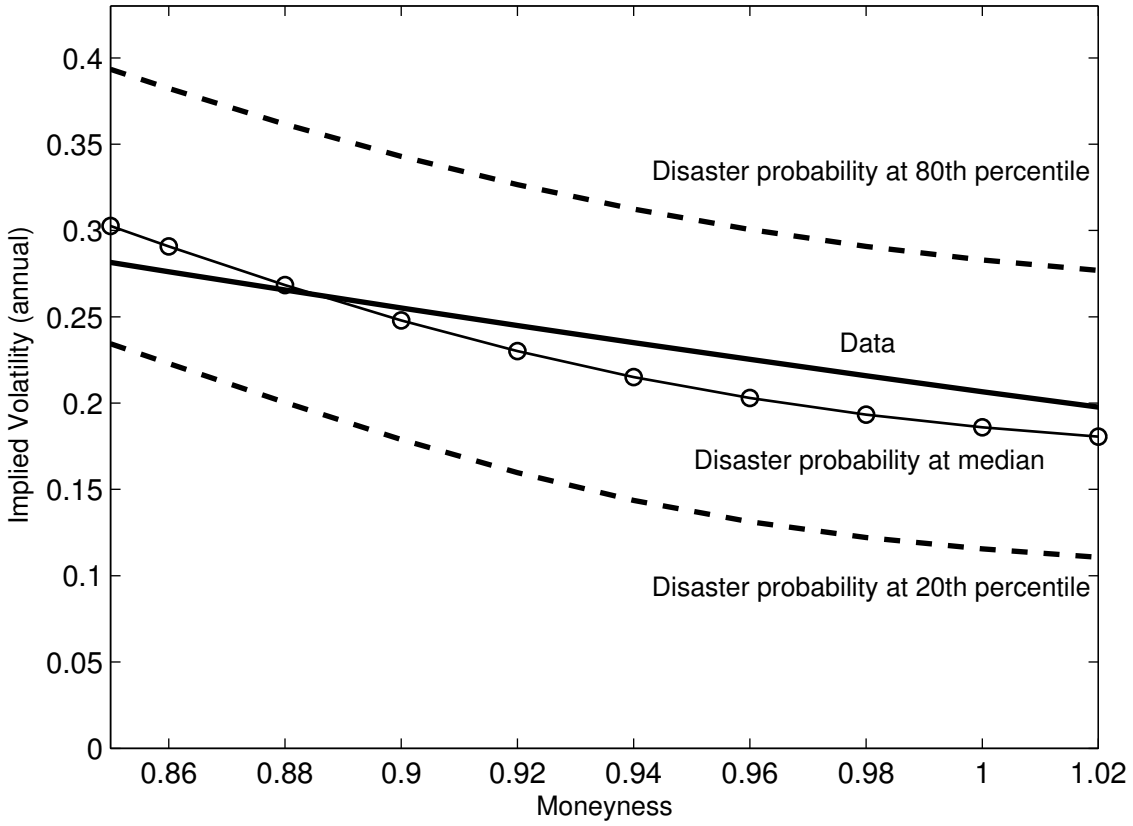

Notes: Implied volatilities on 3-month options as functions of moneyness and of the probability of a disaster for the SDR model.

Figure 10: Fitted values of state variables

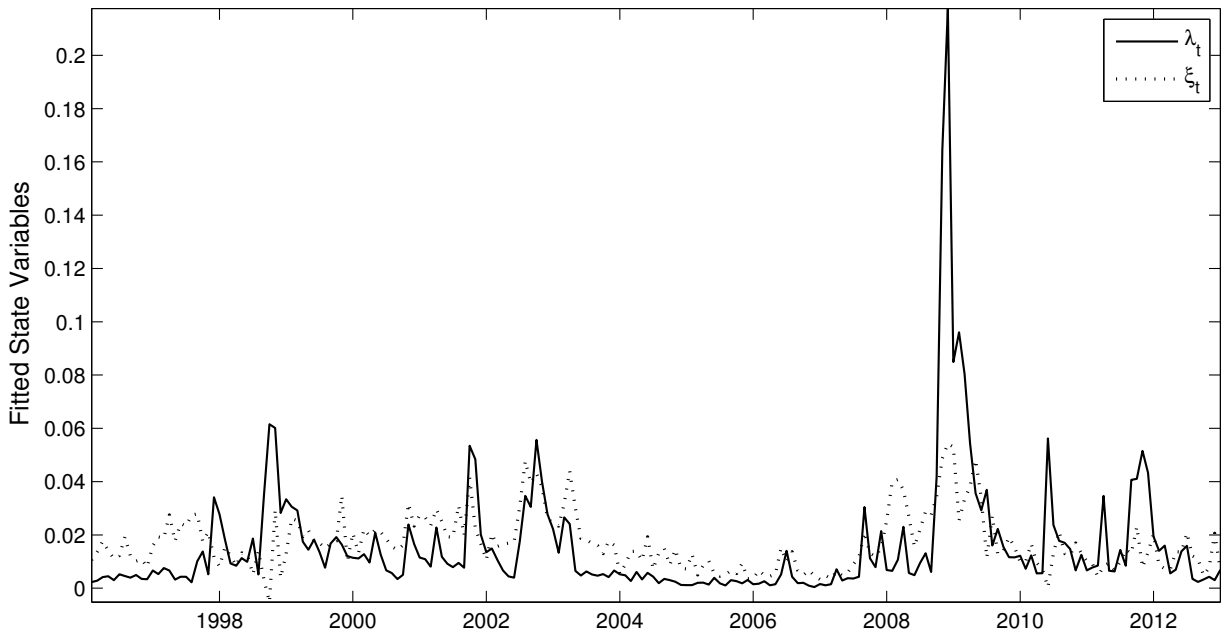

Notes: Fitted state variables in the two-factor SDR model. State variables are fitted to the time series of 1-month implied volatilities of ATM options and OTM options with moneyness equal to 0.85. The variable λ_t represents the (annual) probability of disaster. λ_t reverts to a time-varying value ξ_t.

Figure 11: 3- and 6-month implied volatility time series

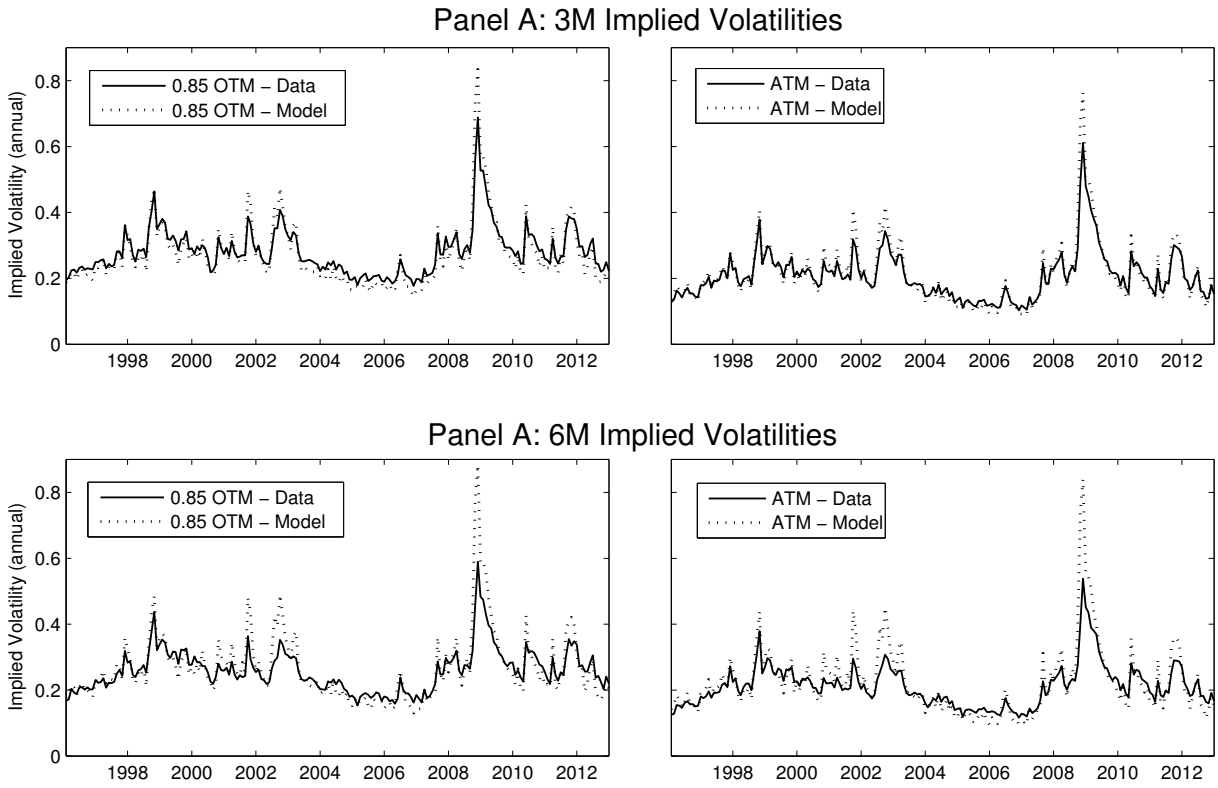

Notes: Monthly time series of implied volatilities on 3- and 6-month options in the data and in the two-factor SDR model. Implied volatilities are calculated for ATM options and OTM options with moneyness equal to 0.85. In the model, we compute option prices using state variables fit to the time series of 1-month ATM and OTM options with moneyness of 0.85.

Figure 12: Implied volatilities as functions of the state in the two-factor SDR model

Panel A: Varying λ_t

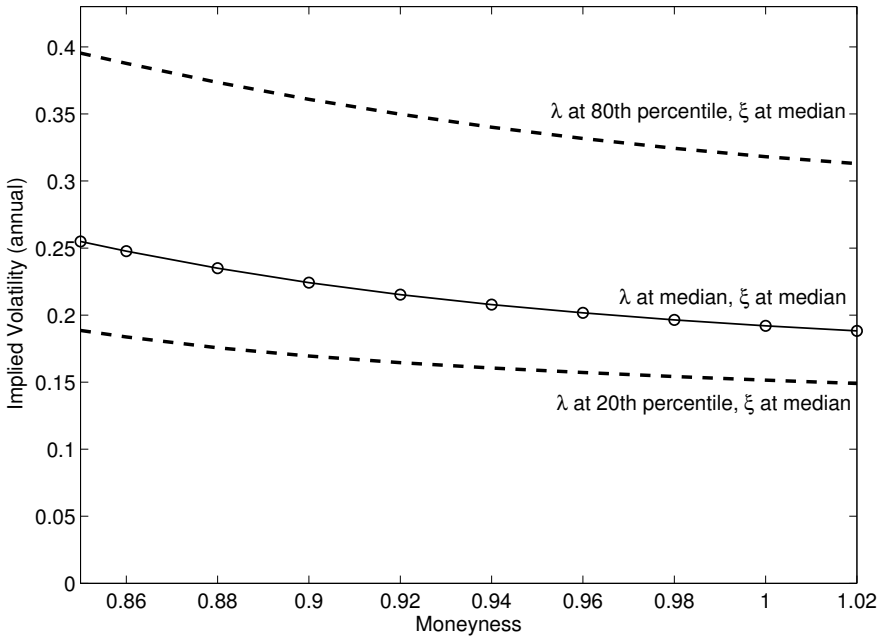

Panel B: Varying ξ_t

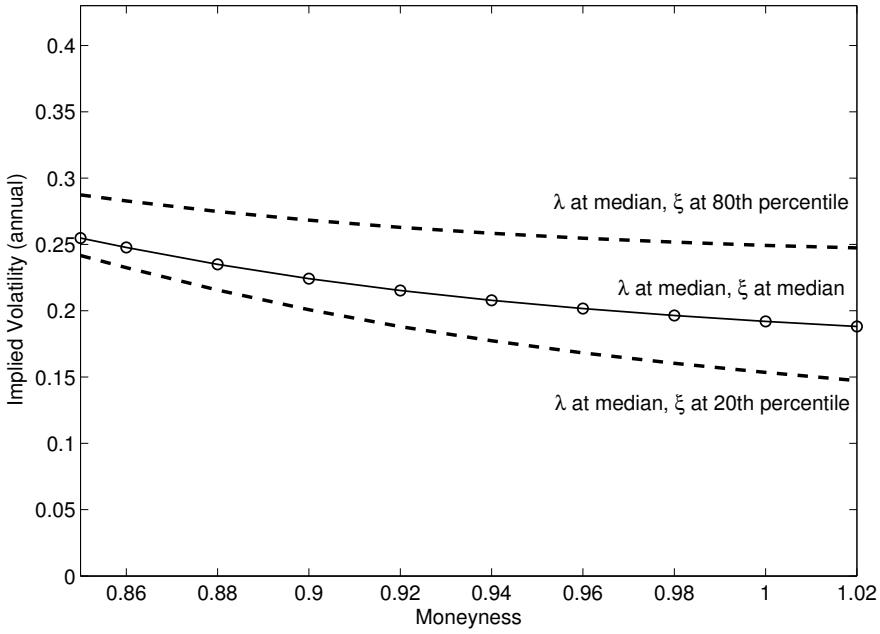

Notes: Implied volatilities on 3-month options as functions of moneyness. The figures show the effects of varying the state variables λ_t (the disaster probability) and ξ_t (the value to which λ_t reverts). Panel A sets ξ_t equal to its median value and varies λ_t, while Panel B sets λ_t equal to its median value and varies ξ_t.

Figure 13: 1-month average implied volatility curve in the data and in the two-factor SDR model

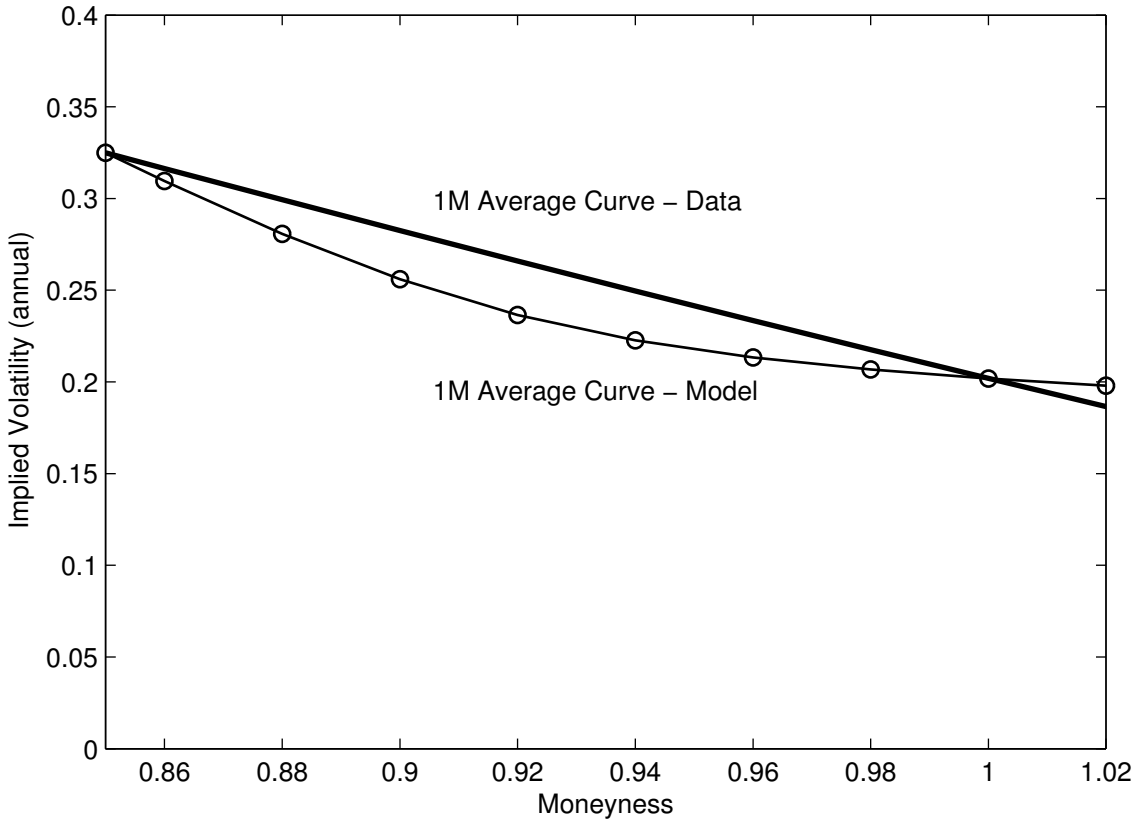

Notes: Average implied volatilities on 1-month options in the data and in the two-factor SDR model. In the model, we compute option prices using state variables fit to the time series of 1-month ATM and OTM options with moneyness of 0.85. We compute an implied volatility for each option price and then take the average.

Figure 14: 3-month average implied volatility curve in the data and the two-factor SDR model

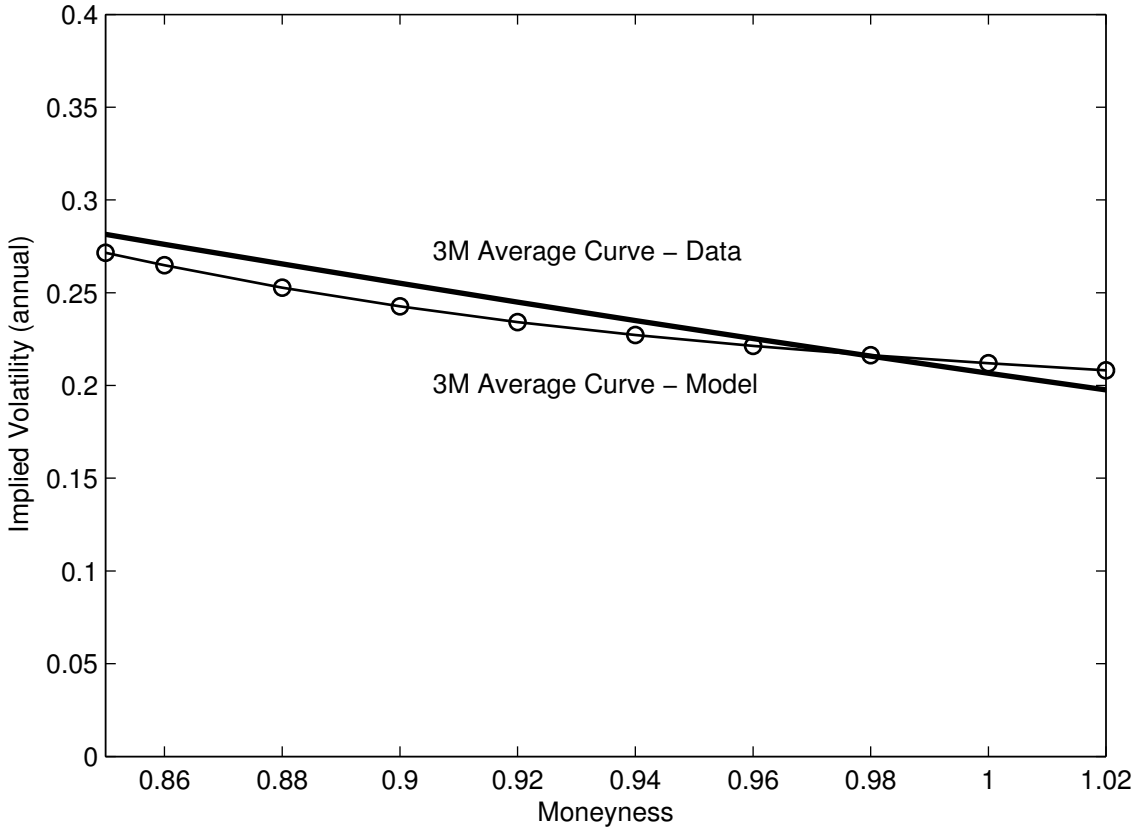

Notes: Average implied volatilities on 3-month options in the data and in the two-factor SDR model. In the model, we compute option prices using state variables fit to the time series of 1-month ATM and OTM options with moneyness of 0.85. We compute an implied volatility for each option price and then take the average.

Figure 15: 6-month average implied volatility curve in the data and in the two-factor SDR model

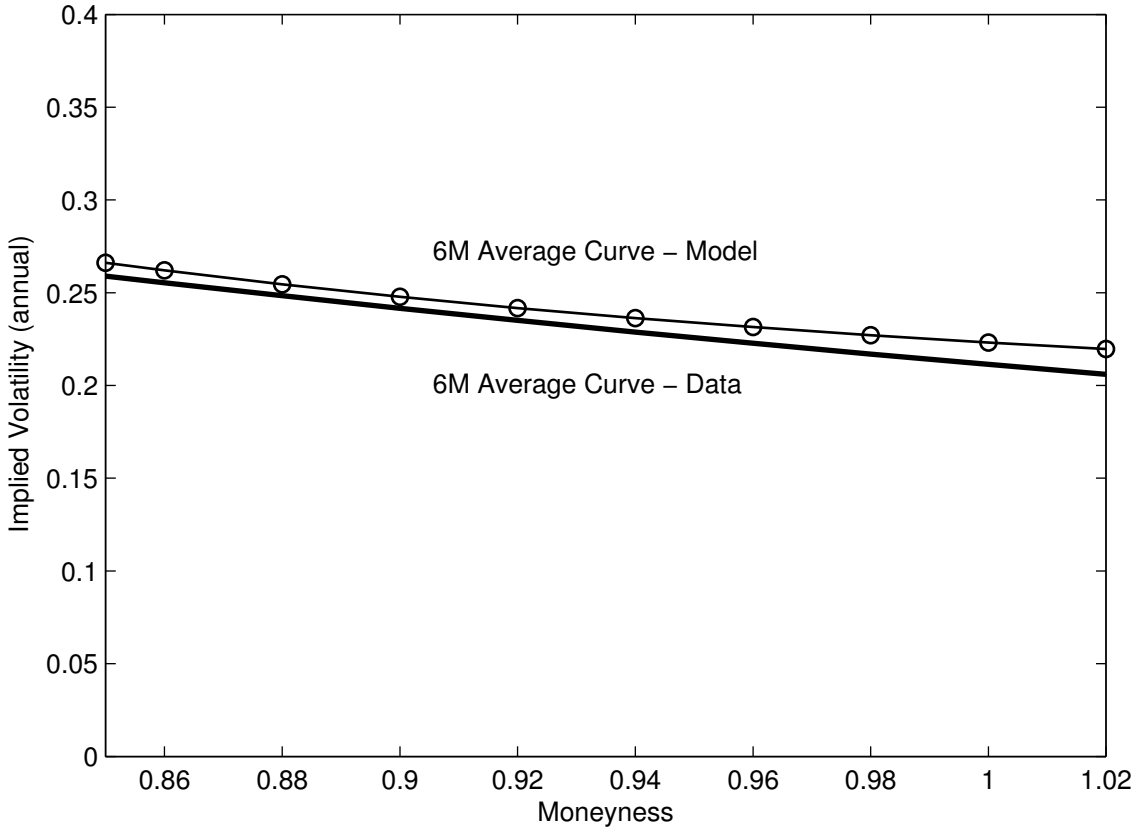

Notes: Average implied volatilities on 6-month options in the data and in the two-factor SDR model. In the model, we compute option prices using state variables fit to the time series of 1-month ATM and OTM options with moneyness of 0.85. We compute an implied volatility for each option price and then take the average.

Figure B.1: Exact versus approximate solution

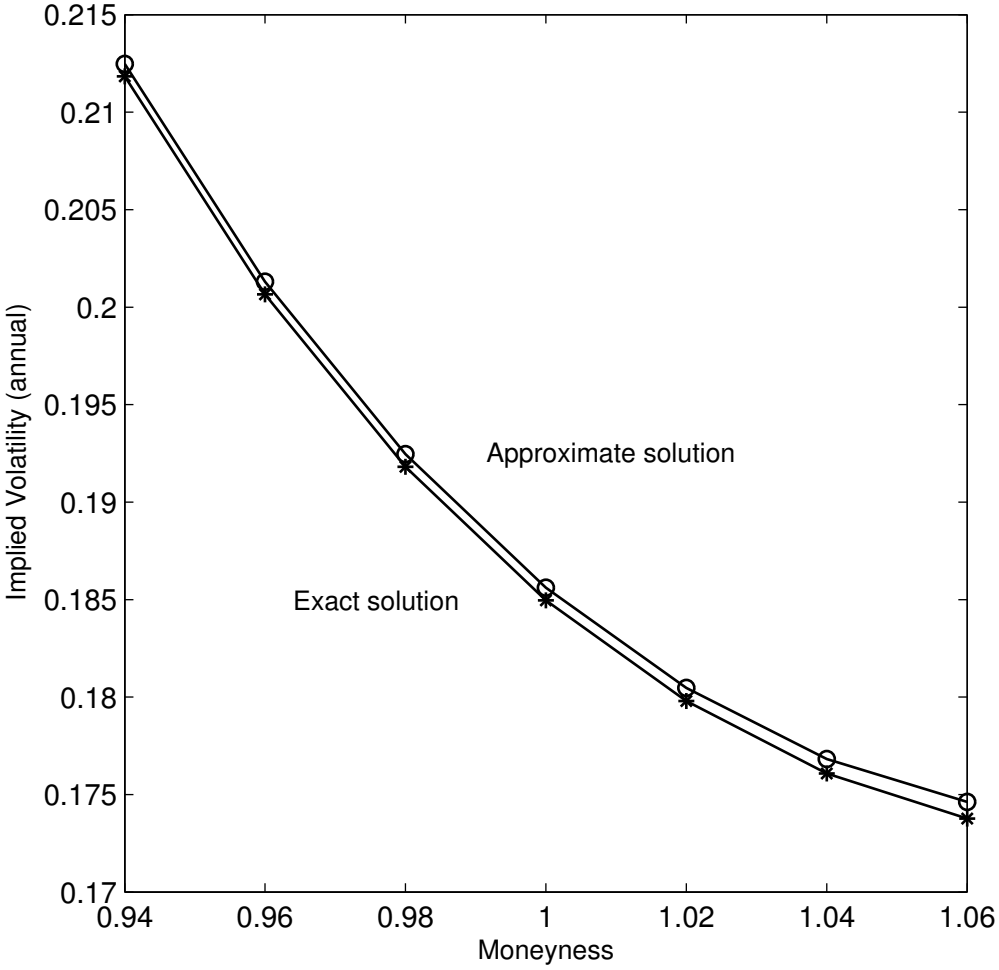

Notes: This figure shows implied volatilities when option prices are computed exactly (line with stars) versus when they are computed using an approximation (line with circles) in the SDR model. Implied volatilities assume that the disaster probability is fixed at its mean and that, in the event of disaster, consumption falls by a fixed amount, namely 30%.

Table 1: Parameter values

	SDR	CDR
Relative risk aversion γ	3.0	5.19
EIS ψ	1	1
Rate of time preference β	0.0120	0.0189
Average growth in consumption (normal times) μ	0.0252	0.0231
Volatility of consumption growth (normal times) σ	0.020	0.010
Leverage ϕ	2.6	5.1429
Average probability of a rare disaster $\bar{\lambda}$	0.0355	0.010
Mean reversion κ	0.080	NA
Volatility parameter σ_λ	0.067	0

Notes: Parameters for stochastic disaster risk (SDR) model and for the benchmark constant disaster risk (CDR) model, in annual terms.

Table 2: Parameter values for the two-factor SDR model

Panel A: λ process	
Mean reversion κ_λ	0.2
Volatility parameter σ_λ	0.1576

Panel B: ξ process	
Mean $\bar{\xi}$	0.02
Mean reversion κ_ξ	0.1
Volatility parameter σ_ξ	0.0606

Panel C: Population statistics of λ	
Median	0.0037
Standard deivation	0.0386
AR(1) coefficient	0.9858

Notes: Parameter values for the two-factor SDR model. The processes are as follows:

$$\begin{aligned} d\lambda_t &= \kappa_\lambda(\xi_t - \lambda_t)dt + \sigma_\lambda\sqrt{\lambda_t}dB_{\lambda,t} \\ d\xi_t &= \kappa_\xi(\bar{\xi} - \xi_t)dt + \sigma_\xi\sqrt{\xi_t}dB_{\xi,t}. \end{aligned}$$

Parameter values are expressed in annual terms. Panel C shows population statistics for the disaster probability λ_t calculated by simulation at a monthly frequency (so that the AR(1) coefficient should be interpreted in monthly terms).

Table 3: Moments for the government bill rate and the market return for the two-factor SDR model

		No-Disaster Simulations			All Simulations			
	Data	0.05	0.50	0.95	0.05	0.50	0.95	Population
$E[R^b]$	1.25	1.68	2.96	3.46	-0.47	2.41	3.37	2.02
$\sigma(R^b)$	2.75	0.34	1.07	2.71	0.48	2.06	7.14	3.69
$E[R^m - R^b]$	7.25	5.40	8.01	12.36	5.30	8.49	14.25	9.00
$\sigma(R^m)$	17.8	13.24	19.26	27.91	14.59	22.59	34.38	24.13
Sharpe Ratio	0.41	0.32	0.42	0.55	0.26	0.39	0.53	0.37
$\exp(E[p-d])$	32.5	28.96	40.63	48.88	22.93	36.95	47.41	35.36
$\sigma(p-d)$	0.43	0.15	0.27	0.47	0.17	0.33	0.59	0.43
AR1$(p-d)$	0.92	0.59	0.79	0.91	0.62	0.82	0.92	0.90

Notes: Data moments are calculated using annual data from 1947 to 2010. Population moments are calculated from simulating data from the two-factor stochastic disaster risk (SDR) model at a monthly frequency for 600,000 years and then aggregating monthly growth rates to an annual frequency. We also simulate 100,000 60-year samples and report the 5th, 50th and 95th percentile for each statistic, both from the full set of simulations and for the subset of samples for which no disasters occur. R^b denotes the government bill return, R^m denotes the return on the aggregate market and $p-d$ denotes the log price-dividend ratio.

Table 4: Moments of state variables in the two-factor SDR model

		No-Disaster Simulations			All Simulations		
	Data	0.05	0.50	0.95	0.05	0.50	0.95
$E[\lambda]$	0.0157	0.0011	0.0083	0.0464	0.0013	0.0117	0.0754
$\sigma(\lambda)$	0.0237	0.0022	0.0118	0.0437	0.0026	0.0153	0.0592
$AR1(\lambda)$	0.7379	0.7604	0.9272	0.9750	0.7825	0.9372	0.9784
$E[\xi]$	0.0170	0.0043	0.0142	0.0429	0.0046	0.0158	0.0495
$\sigma(\xi)$	0.0106	0.0032	0.0083	0.0202	0.0034	0.0090	0.0222
$AR1(\xi)$	0.7035	0.9042	0.9593	0.9837	0.9057	0.9604	0.9842

Notes: Data moments of state variables are calculated using state variables extracted from monthly data on option prices from 1996 to 2012. The parameter λ_t is the probability of a disaster. The probability of a disaster reverts to the state variable ξ_t. We simulate 100,000 17-year samples at a monthly frequency and report the 5th, 50th and 95th percentile for each statistic, both from the full set of simulations and for the subset of samples for which no disasters occur.