Option prices in a model with stochastic disaster risk *

Sang Byung Seo
University of Pennsylvania

Jessica A. Wachter
University of Pennsylvania
and NBER

March 18, 2013

Abstract

In a challenge to models that link the equity premium to rare disasters, Backus, Chernov, and Martin (2011) show that data on options imply negative events that are far smaller than these models suggest. We show that this result depends critically on the assumption that the probability of the rare event is constant. That is, a model with stochastic jumps in consumption can simultaneously explain options data and the equity premium. Indeed, such a model delivers an excellent fit to implied volatilities, despite being calibrated to match the equity premium and equity volatility alone.

*First draft: March 18, 2013. Seo: sangseo@wharton.upenn.edu; Wachter: jwachter@wharton.upenn.edu.
We thank Domenico Cuoco, Mikhail Chernov, Ivan Shaliastovich, and seminar participants at the Wharton school for helpful comments.

Option prices in a model with stochastic disaster risk

In a challenge to models that link the equity premium to rare disasters, Backus, Chernov, and Martin (2011) show that data on options imply negative events that are far smaller than these models suggest. We show that this result depends critically on the assumption that the probability of the rare event is constant. That is, a model with stochastic jumps in consumption can simultaneously explain options data and the equity premium. Indeed, such a model delivers an excellent fit to implied volatilities, despite being calibrated to match the equity premium and equity volatility alone.

1 Introduction

Century-long evidence indicates an economically significant equity premium: an expected return from holding equities over short-term government 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 rare events. These models include Rietz (1988), Barro (2006) and Weitzman (2007). In these models, consumption 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 rare events. 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 unusual 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 offer a second, largely independent source of information. It is therefore of interest to see what option prices can teach us about the probability of rare events.

Recent work seeks to address exactly these questions. Backus, Chernov, and Martin

[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).

(2011) derive option prices in a rare events model similar to that of Rietz (1988) and Barro (2006). They find 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. This result suggests that these rare events are not the source of risk behind the equity premium. 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.[3]

In this paper, we follow the approach of Backus, Chernov, and Martin (2011) but with an important exception. Like the earlier literature discussed above, Backus et al. assume that the arrival of rare events is iid. 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 seek to incorporate dynamics that can also account for such volatility (Gabaix (2012), Gourio (2011), Wachter (2012)). In this paper, we derive option prices in the model of Wachter. Importantly, we use the same parameters that have been shown to explain equity returns in that paper. We show that allowing for a stochastic probability 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.

[3]Backus, Chernov, and Martin (2011) derive the distribution of consumption growth implied by option prices. Besides using this to determine the probability of rare events implied by options, they also suggest such a calibration, which has more frequent and less severe events, is a useful alternative to that of Rietz (1988) and Barro (2006) in that it can also explain the equity premium. However, because of the frequency of the negative events, this calibration is unable to explain the exceptionally low volatility enjoyed by the U.S. during the 60-year post-war period. A period with such low volatility would have less than a 1 in 100,000 chance of occurring.

As we show, the type of model we consider nests that considered by Backus, Chernov, and Martin (2011). This allows us to break apart the differences between the models, and consider what exactly is causing the difference. We conclude that it is indeed the time-variation in the probability of disaster that accounts for the majority of the difference. The reasoning is two-fold. First, by raising volatility during normal times, 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 (very low) volatility of consumption growth. Second, 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 otherwise be.

Our findings relate to those of Gabaix (2012), who also reports implied volatilities in a model with rare events. The papers differ in several ways. First, for tractability, Gabaix derives option prices for under different assumptions than those used to price equities. In our work, the underlying assumptions are the same for the two asset classes. Second, because our model nests that of Backus, Chernov, and Martin (2011), we can drill down to uncover the explanation for why we find something different. Finally, our underlying model is different from those used in Gabaix. He assumes a linearity-generating process (Gabaix (2008)) for prices. In his calibration, the investor has power utility and the sensitivity of dividends to changes in consumption is varying; however, the probability of a disaster is not. In our calibration, the investor has recursive utility and the probability is time-varying. We argue below that this has important consequences for the model's ability to account for implied volatilities.

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 indicate the difficulties of matching implied volatilities assuming either only external habit formation, or a constant probability of a disaster. Also related is Drechsler and Yaron (2011), who focus on the volatility premium and its predictive properties, rather than on implied volatilities. Besides the difference in focus, the papers also differ in the economic mechanism. 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 second related strand of research focuses on uncertainty aversion or exogenous changes in confidence. Like some of the work mentioned above, Drechsler (2012) builds on Bansal and Yaron (2004), but incorporates in dynamic uncertainty aversion (Knightian uncertainty). 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 asset returns Section 3 discusses the fit to the data and Section 4 concludes.

2 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 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.[4] Duffie and Skiadas (1994) characterize the state-price density as

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

There is an equilibrium relation between utility V_t, consumption C_t and the disaster probability λ_t. Wachter (2012) shows that

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

[4]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).

where a and b are constants given by

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

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

It follows that

$$\pi_t = \exp\left(\eta t - \beta b\int_0^t \lambda_s ds\right)\beta^\gamma C_t^{-\gamma}e^{a+b\lambda_t}, \tag{7}$$

where

$$\eta = \beta(1-\gamma)\log\beta - \beta a - \beta.$$

Details are provided in Appendix A.1.

Following Backus, Chernov, and Martin (2011) and Wachter (2012), we assume a simple relation between dividends and consumption: $D_t = C_t^\phi$, for leverage parameter ϕ. 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

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

Wachter (2012) shows that the stock price can be written explicitly as

$$F(D_t, \lambda_t) = D_t G(\lambda_t), \tag{8}$$

where the price-dividend ratio G is given by

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

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

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

where

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

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 $t+\tau$. Because the payoff on this option at expiration is $(K - F_{t+\tau})^+$, it follows from the absence of arbitrage that

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

At each time t, options are issued at a range of strike prices K near $K = F_t$ (namely, at the money). Let $K^n = K/F_t$, the normalized strike price (or "moneyness"), and define

$$P^n(\lambda_t, T-t; K^n) = E_t\left[\frac{\pi_T}{\pi_t}\left(K^n - \frac{F_T}{F_t}\right)^+\right]. \tag{9}$$

We will establish below that P^n is a function of λ_t and the time to expiration alone. Clearly $P^n = P/F_t$. 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^n.[5]

[5] 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

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

where

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

It follows from (8) that

$$\frac{F_T}{F_t} = \frac{D_T}{D_t}\frac{G(\lambda_T)}{G(\lambda_t)} \tag{10}$$

Moreover, it follows from (7) that

$$\frac{\pi_T}{\pi_t} = \left(\frac{C_T}{C_t}\right)^{-\gamma} \exp\left\{\int_t^T (\eta - \beta b \lambda_s)\, ds + b(\lambda_T - \lambda_t)\right\}. \tag{11}$$

Because the distribution of consumption and dividend growth between t and T, and because the distribution of λ_s for $s = t, \dots, T$ depends only on λ_t, normalized put prices (and therefore implied volatilities) are a function of λ_t and the horizon alone.

We therefore calculate (9), assuming the parameter values and distribution of disaster sizes in Wachter (2012). We first approximate the price-dividend ratio $G(\lambda_t)$ by a log-linear function of λ_t, as described in Appendix B.1. 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)).[6]

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

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

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

[6] At the parameter values we consider, the Feller condition (Feller (1951)), namely $2\kappa\bar{\lambda} > \sigma_\lambda^2$, is satisfied implying a finite density at zero.

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 power utility with a constant probability of log-normal jumps in consumption. They call this the consumption-based model, to distinguish it from the options-based model which is reduced-form and designed to fit options data. There are two apparent differences between their consumption-based model model and the CDR model. First, Backus et al. assume power utility. Second, their model assumes discrete time while the present model assumes time is continuous. However, neither difference turns out to matter.

First, though Backus, Chernov, and Martin (2011) 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). In terms of the continuous/discrete-time distinction, Appendix C shows that a calibration of the CDR model when observed in discrete time is identical to the model assumed by Backus et al. (based on that of Merton (1976))

There is one further difference in implementation between the present paper and that of Backus, Chernov, and Martin (2011). Backus et al. define option payoffs not in terms of the price, but in terms of the total return. Moreover, it is not possible to compute option payoffs defined in terms of prices at their parameters because these prices do not converge.[7] To evaluate the possible quantitative importance of this difference, we calculate implied volatilities in a calibration similar to that of Backus et al., but where we set the riskfree rate to be a high value so prices do converge. Because the dividend yield is low, this calibration will not artificially understate the size of the of replacing changes in prices with returns. The results are reported in Appendix D. While there is a difference in implied volatilities, it is confined to in-the-money options. We conclude that the distinction between options on returns and prices is not an important determinant of the difference between the results they report and ours. However, it does affect the computation method used. When reporting results for their calibration, we use the method outlined in their paper. We have verified that the two methods produce the same answer (given the appropriate dividend yield adjustment) in settings where they both apply.

3 Comparison with the data

3.1 Data

Our sample consists of daily data on implied volatilities, options prices, volume and open interest for European put options on the S&P 500 index from OptionMetrics. We extract

[7]Backus, Chernov, and Martin (2011) do not make a direct assumption about the discount rate. However, they do assume that the riskfree rate is 2%; this assumption is used to calculate implied volatilities. For consistency, we maintain this assumption when using their calibration, and back out the discount rate assuming recursive utility with EIS equal to one. The assumption of a 2% riskfree rate implies, given their other parameters, that prices will not converge.

observations with maturity of three months (91 days) from daily options data between 1996 and 2011. 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 0.5 (Figlewski (2010)) and those with zero volume and open interest smaller than one hundred contracts (Shaliastovich (2009)). Note that the available range of liquid options will not be constant for each trade day, leading to potential bias. We therefore follow standard practice in using cubic spline interpolation and flat extrapolation (e.g. Carr and Wu (2009)) to infer volatilities over the relevant range.

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 (2012)), the two models are in a sense equivalent, and the comparison is valid. Second, this comparison will also allow us to reconcile our results with theirs.[8] However, in 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. Some of these differences are less important

[8]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)).

than what may first appear: 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. Backus, Chernov, and Martin (2011) assume 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 the strike price, 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, the implied volatilities for the in-the-money options implied by the consumption-based model are less than 10%, far below the option-based implied volatilities, which are over 20%.

In contrast, implied volatilities for the stochastic disaster risk (SDR) model are much closer to the data, both in terms of the level and the slope. For out-of-the-money options, the implied volatilities are about 24%, close to the data value of 21%. Their is a downward slope, just as in the data, but it is much smaller than that of the CDR model. In the money options have implied volatilities of about 22%. The fit to the data is not perfect, but it is quite good, given that the model parameters were chosen without regard to the options data. Moreover, the data themselves contain some noise, as the pattern in implied volatilities exhibits substantial time variation (e.g. Benzoni, Collin-Dufresne, and Goldstein (2011)).

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 effect 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. The figure also shows that the shape of the curve depends on the probability, when the curve is at the 20th percentile, it is noticeably flatter than when it is at the 80th percentile. Most importantly, the curve at the median value also lies near to that in the data.

3.4 Discussion

What is it about the SDR model that enables it to match the data? 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 the increased probabilities of very bad consumption realizations in the consumption-based model, as opposed to an implied consumption distribution from the option-based model. Therefore, this seems like an appropriate place to start. While the consumption distribution conditional on a disaster is actually more extreme in the case of the SDR model the leverage parameter used by Backus et al. is much higher than in the SDR model, 5.1 versus 2.6, implying that stock returns have fatter tails. Option prices will depend on both the consumption process and the dividend process, since the consumption process determines the pricing kernel. It is reasonable, therefore, to attribute the difference in the implied volatilities to the difference in the leverage parameter. Indeed, Gabaix (2012) conjectures that this is why Backus et al. find the results that they do.

Figure 6 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. The main difference between the calibration is that, for lower leverage, the implied volatility curve shifts downward. Changing leverage does not only change the disaster distribution, it also changes normal-times volatility. This is what accounts for the large difference for the in-the-money options. However, the fact that the graph does not flatten for the out-of-the money options suggests that it is not the severity of the disasters that drives the differences in the results.

If not the severity of the disasters, then what does drive the difference? We consider another aspect of the CDR calibration, namely, normal-times volatility. In our benchmark case, 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%. 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 6, the result is shown in the line denoted "higher normal-times volatility."

As Figure 6 shows, increasing the normal-times volatility by even a small amount has a dramatic effect on implied volatilities (note that raising normal-times volatility while keeping the other parameters unchanged does imply that the resulting stock return volatility will exceed that in the data). The implied volatility curve is 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. Thus the difference in normal-times volatility could indeed be driving part of the difference between the CDR and the SDR models. 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.

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 5. 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 surprisingly 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, emphasized by Wachter (2012), is high stock 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 lead to an upward slope for out-of-the money put options, high normal-period volatility imply 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.

It is not obvious that this should be the case. Part of the puzzle of the implied volatility smile is that implied volatilities are higher than realized volatilities. The model can fit this aspect of the data because of a mechanism 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 (7) with one in which $b = 0$.[9] Because b determines the risk premium due to covariance with λ, 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 are surprising given that incorporating stochastic moments has not, in pre-

[9]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

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

.

vious literature, appeared to affect the shape of the smile. Indeed, Backus, Chernov, and Martin (2011) reasonably make this very argument:

> 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.

How is it, then, that this paper comes to such a different conclusion? The reason may arise from the fact that the literature referred to in the previous quotation has 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, stock market volatility arises 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 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 (at least in a model with recursive preferences) as well as generates the excess volatility observed in stock prices.

The model that we have developed in this paper is deliberately simple and parsimonious: we show that a model that can match equity prices with no modifications can also match option prices. However, we anticipate that there are a number of interesting features of option and stock return data that cannot be matched using a model with a single state variable that follows a square root process. We look forward to exploring generalizations, guided by both the physical and risk neutral distributions of returns, in 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 (11), we require both to be expressed in terms of C_t and λ_t. We derive the result for the stochastic term first.

It follows from (4) that

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

Wachter (2012) shows that continuation utility V_t can be expressed in terms of C_t as follows:

$$V_t = J(\beta^{-1} C_t, \lambda_t), \tag{A.3}$$

where

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

and

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

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

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

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

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

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

Finally, we collect constant terms:

$$\eta = \beta(1-\gamma)\log\beta - \beta a - \beta \tag{A.10}$$

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

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

A.2 The iid limit of the SDR model.

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

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

Applying L'Hopital's rule implies that

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

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

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

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

We consider the limit of π_t/π_0 as σ_λ approaches zero. Note that for computing asset prices, we only care about ratios of the state price density at different points of time, and so it suffices to compute this quantity. It follows from (A.1), (A.8) and (A.9) that

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

which is equivalent to the result one obtains by calculating the state price density in the iid case when the EIS is equal to 1. Note that this result is not automatic, but rather holds only if we choose lower of the two roots of (A.7) as we do in this paper, a point made in a related context by Tauchen (2005).

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.

Let $\pi_{p,t}$ be the state price density assuming power utility with discount rage 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 (2012) 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 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 $\exp\left\{a_\phi(\tau) + b_\phi(\tau)\lambda^*\right\}$. With this in mind, we define the notation

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

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

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

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 (9) 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 (10), (11), 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), (6) and (B.3), respectively. Then (B.5) can be rewritten as

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

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

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

and let

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

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

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

and therefore

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

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 The Poisson-normal model in continuous time

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

$$\frac{dC_t}{C_{t^-}} = \mu dt + \sigma dB_t + (e^{Z_t} - 1)dN_t$$
$$Z_t \sim N(\theta, \delta^2), \quad N_t \sim \text{Poisson}(\bar{\lambda})$$

Ito's Lemma implies

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

and therefore,

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

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

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

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

D Options on total return versus price

In Figure D.2, we compare implied volatilities calculated from options written on the return of a security and from options written on the price of the security. We use parameter values as in the benchmark CDR case, however, so that the price-dividend ratio converges, we consider a higher riskfree rate, namely 8%. For out-of-the-money options, the two smiles are nearly identical. For in-the-money options, they diverge, with the curve corresponding to options written on the price curving upwards (becoming more "smile-like"). We conclude that, for out-of-the-money options, the issue of return vs. price is not quantitatively important.

References

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

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long-run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866.

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

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.

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, 2009, Variance Risk Premiums, *Review of Financial Studies* 22, 1311–1341.

Chacko, George, and Luis Viceira, 2005, Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets, *Review of Financial Studies* 18, 1369–1402.

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.

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.

Feller, William, 1951, Two Singular Diffusion Problems, *The Annals of Mathematics* 54, 173–182.

Figlewski, Stephen, 2010, Estimating the implied risk neutral density for the US market portfolio, in Tim Bollerslev, Jeffrey Russell, and Mark Watson, eds.: *Volatility and Time Series Econometrics: Essays in Honor of Robert Engle* (Oxford University Press, Oxford, UK).

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

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.

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.

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.

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.

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

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.

Wachter, Jessica A., 2012, Can time-varying risk of rare disasters explain aggregate stock market volatility?, forthcoming, *Journal of Finance.*

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.

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

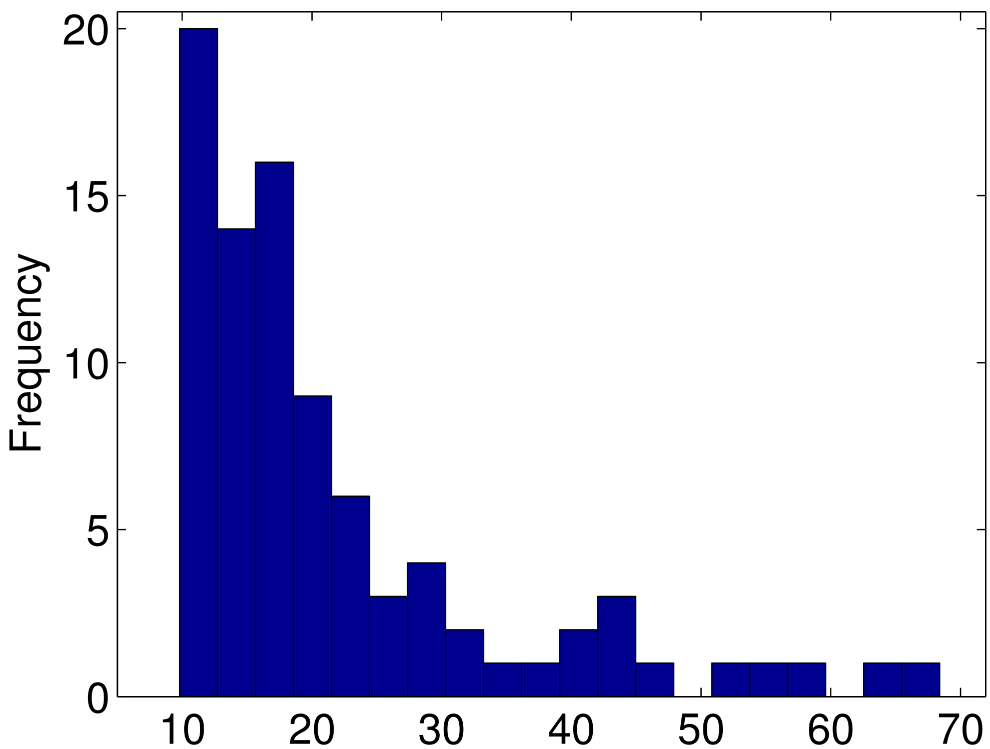

Notes: Histogram of large consumption declines (in percentages). The distribution is calculated by Barro and Ursua (2008) based on a century-long dataset of 22 countries. The histogram is for the quantity $1 - e^{Z}$ in the model.

Figure 2: Probability density functions for consumption declines

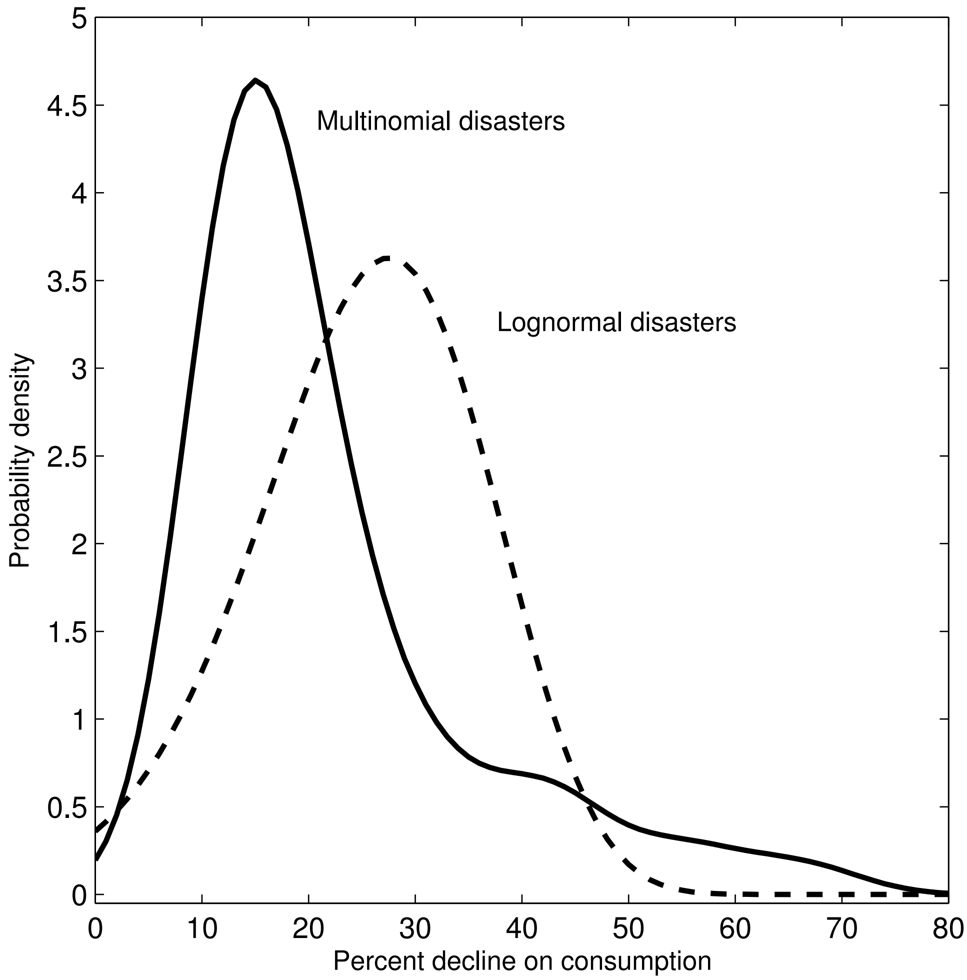

Notes: The probability density functions (pdfs) for consumption declines for log-normally distributed disasters and for the multinomial distribution assumed in the stochastic disaster risk (SDR) model. In the case of the SDR model, the pdf approximates the multinomial distribution from Barro and Ursua (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: Implied volatilities in the SDR and CDR models

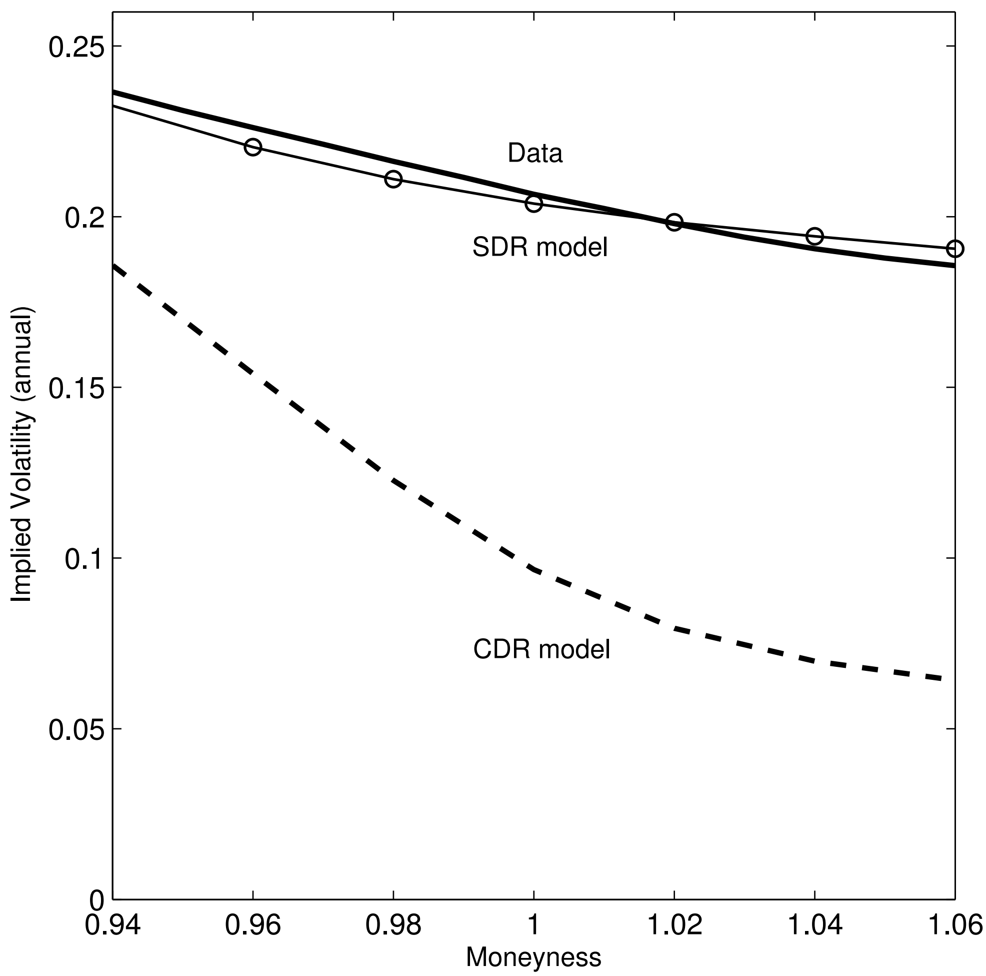

Notes: The figure shows implied volatilities for three-month options as a function of moneyness for the stochastic disaster risk (SDR) model, for the constant disaster risk (CDR) model (under the benchmark calibration given in Table 1) and in the data.

Figure 4: Conditional implied volatilities in the SDR model

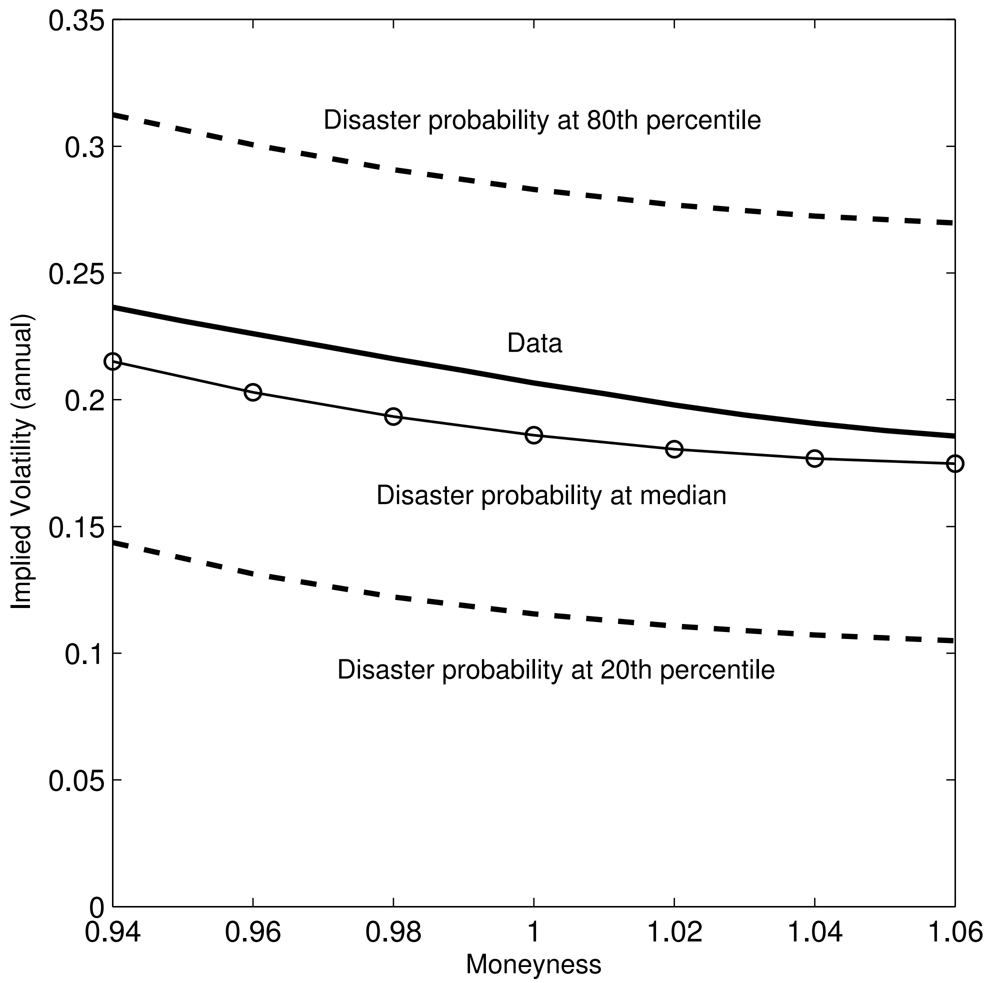

Notes: This figure shows implied volatilities at different values of the state variable, the probability of a disaster, in the stochastic disaster risk model.

Figure 5: Implied volatilities in the SDR model and in an alternative calibration of the CDR model

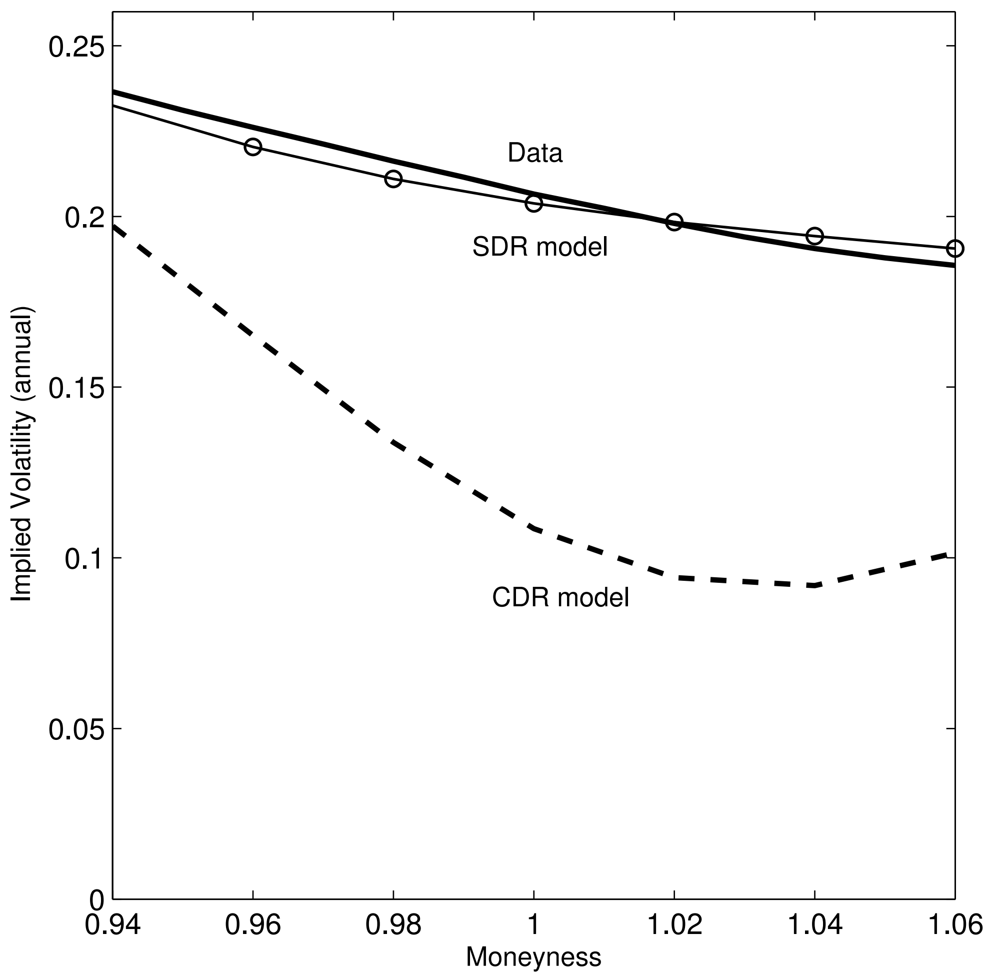

Notes: The figure shows implied volatilities for three-month options as a function of moneyness for the stochastic disaster risk (SDR) model, for the constant disaster risk (CDR) model and in the data. The parameters for the CDR model are chosen to be identical to those of the SDR model, except that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

Figure 6: Comparative statics for the constant disaster risk model

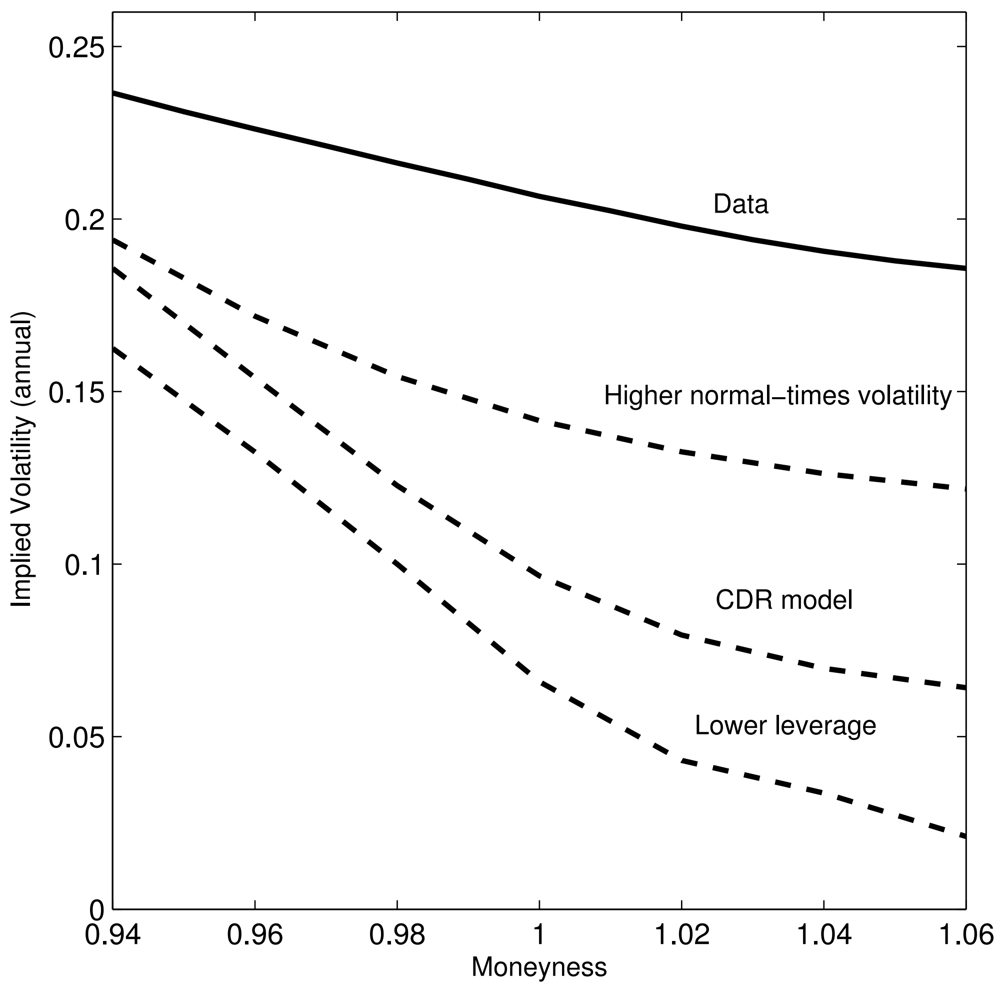

Notes: This figure shows implied volatilities for three-month options as a function of moneyness in the data and for three parameterization of the constant disaster risk model. The line labeled "CDR" shows the benchmark calibration. The line labeled "higher normal-times volatility" raises the volatility 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 7: Evaluating the role of recursive utility

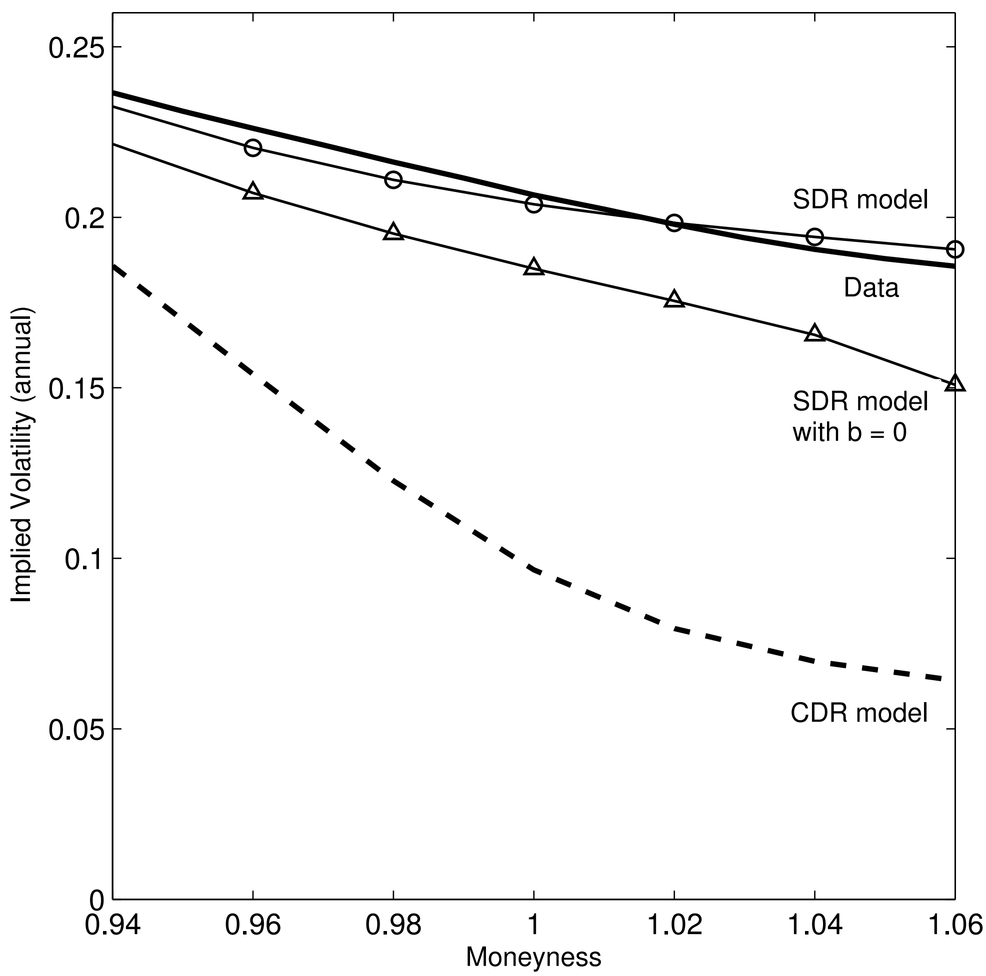

Notes: This figure shows implied volatilities for three-month options as a function of moneyness in the data, in the CDR model (calibrated as in Table 1) and in the SDR model. Also shown are option prices in the SDR model computed under the assumption that the premium associated with time-variation in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

Figure B.1: Exact versus approximate solution

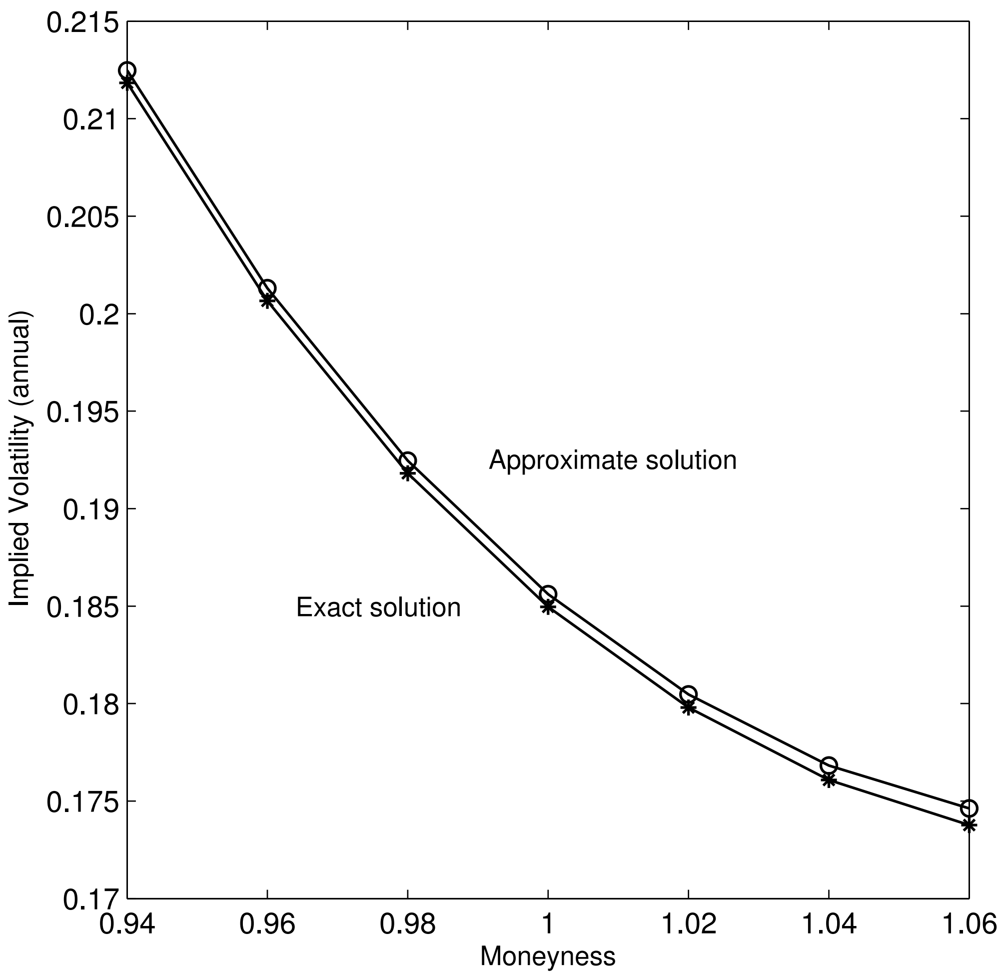

Notes: This figure shows implied volatilities when option prices are computed exactly (line with stars) versus when they are computed using an approximation (line with circles) in the SDR model. Implied volatilities assume that the disaster probability is fixed at its mean and that, in the event of disaster,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

Figure D.2: Options on returns versus prices

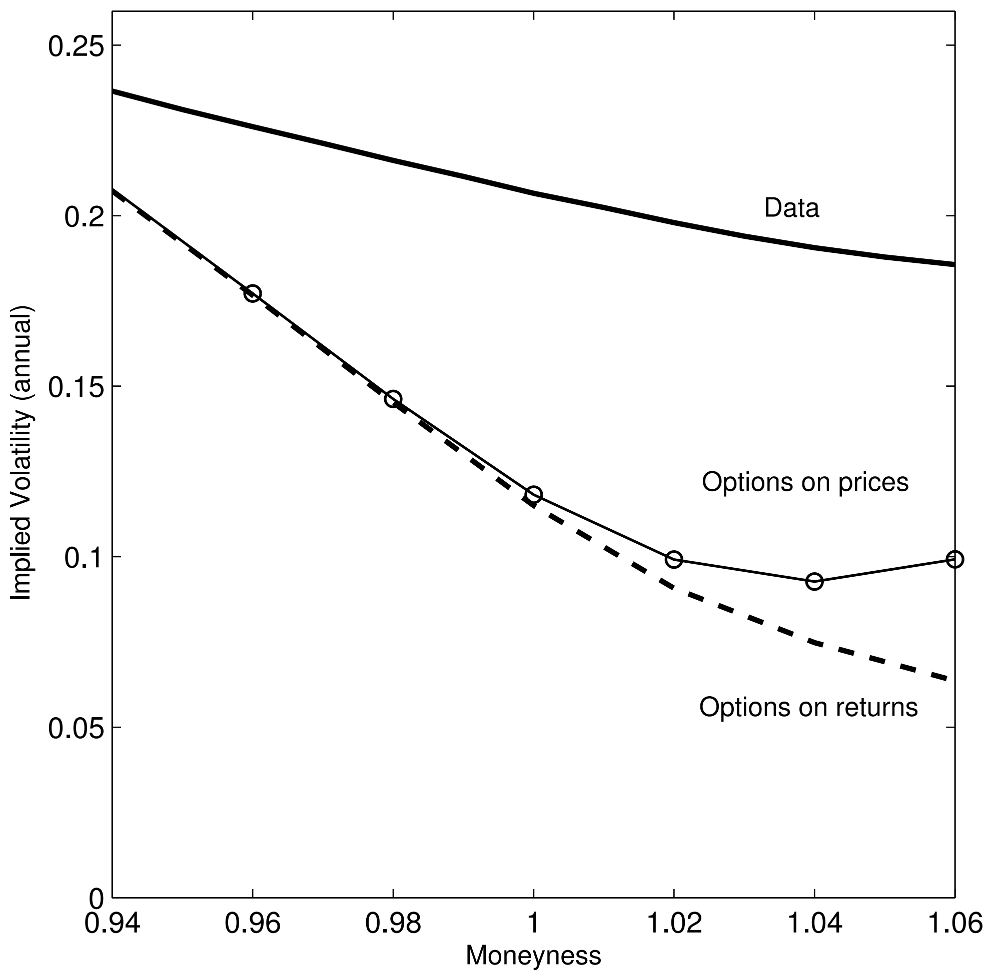

Notes: This figure shows implied volatility curves for 3-month options as a function of moneyness for the benchmark CDR model, except that we adjust time preference parameter β to match 8% risk free rate. The dashed line shows options on returns while the line with circles shows options on prices.

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.