Internet Appendix for

"Uncertainty, Time-Varying Fear, and Asset Prices" *

Itamar Drechsler

This Internet Appendix serves as a companion to the paper "Uncertainty, Time-Varying Fear, and Asset Prices". It reports results not reported in the main text due to space constraints. I present results in the order they appear in the main text.

I. Discussion of Alternative Approaches

I discuss potential alternative approaches to an equilibrium model aimed at jointly capturing the properties of equity returns, option prices, and the variance premium.

Given their popularity in the asset-pricing literature, it is natural to consider models based on habits (e.g., Campbell and Cochrane (1999), Menzly, Santos, and Veronesi (2004)) as a potential alternative equilibrium approach to generating a large equity return and index option price premium. In habits models, variation in the surplus consumption ratio, or equivalently, risk aversion, drives variation in the conditional moments of prices, such their risk premiums and conditional variance. Having this single source drive all variation in valuations and risk premiums creates some challenges for simultaneously confronting the targets considered by my paper. First, since shocks to risk aversion under habits are driven by innovations in i.i.d consumption, they are relatively

*Citation format: Drechsler, Itamar, 2012, Internet Appendix to "Uncertainty, Time-Varying Fear, and Asset Prices", *Journal of Finance* [vol], [pages], http://www.ajajof.org/IA/[year].asp. Please note: Wiley-Blackwell is not responsible for the content or functionality of any supporting information supplied by the authors. Any queries (other than missing material) should be directed to the authors of the article.

‘smooth’. In addition, risk-aversion is typically modeled as quite persistent to match the persistence in the price-dividend ratio. However, matching options prices, the variance premium, and conditional variance dynamics requires some abrupt, non-normal shocks and relatively faster mean reversion. Hence, capturing the variation in these series along with the price-dividend ratio using only the surplus consumption ratio represents a serious challenge to this approach. Second, to get equity return realizations that are large and carry a high price of risk, habits models will generally require large movements in consumption. This is problematic for habits models that follow Campbell and Cochrane (1999) in specifying an i.i.d. endowment process. Du (2010) shows that the implied-volatility skew implied by the calibrations of Campbell and Cochrane (1999) and Menzly, Santos, and Veronesi (2004) are very off far from the data and notes that alternative calibrations of these models cannot resolve this inconsistency.

Another popular paradigm in the asset pricing literature is the rare disasters framework of Rietz (1988) and Barro (2006) and its extensions to time-varying disaster intensity by Gabaix (2011) and Wachter (2011). These models embed a ‘peso problem’, whereby investors price in the possibility of a large negative drop in consumption even though one does not appear in US historical data. This allows the model to match the large equity premium with low levels of risk aversion. In principle such a model should be capable of producing a steep implied volatility skew and a large variance premium, since consumption disasters generate large realized equity returns that are highly priced. However, Backus, Chernov, and Martin (2011) argue that disasters models calibrated as in Barro (2006) generate an implied volatility skew that is actually far too steep relative to the data. In addition, to generate the empirically observed conditional variance dynamics and higher moments of realized returns, such models would require a lot of variation in the conditional disaster intensity, and this is likely to worsen their counter-factual implications for the implied volatility surface. The model in this paper does not rely on a ‘peso problem’. The benchmark model used by the agent fits the observed data well and the alternative models about which the agent is concerned are statistically difficult to distinguishable from it.

Du (2010) combines both of the above mechanisms to generate the implied-volatility skew. He embeds rare consumptions disasters into a model with external habits, allowing him to address some of the drawbacks of habits for option pricing. The rare disasters in his model induce a large, negative return realization that carries a high price of risk. Due to habits, larger jumps are realized when risk aversion is already high, so jumps carry a large risk premium and produce a high value for out-of-the-money (otm) puts. Moreover, this allows his model to use smaller disasters than in Barro (2006) and avoid the excessively steep implied volatility skew it implies. While disasters help the model in Du (2010) to address these challenges, there remain some important limitations. First, the model equates aggregate consumption with dividends, and equity is modeled as the claim to aggregate consumption. While this generates higher risk prices for 'dividends', this comes at the cost of matching the properties of dividends. More importantly perhaps, this implies that the model must generate excessively volatile changes in the price-dividend ratio in order to match unconditional return volatility. In addition, while the model does generate a high unconditional return volatility, it does not match the dynamic properties of conditional variance or the higher moments of equity returns. It also does not match the properties of the variance premium, an important direct measure of the option price premium. Unlike disasters models, which use a low risk aversion and hence 'resolve' the equity premium puzzle, the combined habits-disasters model still requires a high average risk aversion of 34 despite the embedded 'peso problem'. In contrast, the model in this paper uses a risk aversion of only 5.

Benzoni, Collin-Dufresne, and Goldstein (2011) construct a model where there is a rare and large downward jump in a persistent component of consumption growth. In their model, the representative agent learns about which of two regimes controls the intensity of this rare jump. Following Bansal and Yaron (2004), this representative agent has Epstein-Zin (1989) preferences with a preference for early resolution of uncertainty and an IES greater than one. This implies that a large negative shock to growth causes both a big drop in stock prices and an increase in marginal utility. The risk of such a shock increases the prices of otm put options and leads to a steep implied volatility skew. BCDG focus on the stock market crash of 1987, viewing it as being caused by an update in the representative agent's perceived likelihood of being in the high jump intensity regime. The increased perception of being in the 'bad' regime causes a stock market crash and a persistent steepening in the implied volatility skew, though the jump realization that catalyzes this update is relatively small. Similar to the model of this paper, and in contrast to the consumption disasters models, jump shocks in the BCDG model hit the rate of cash flow *growth* rather than its *level*, so the levels of consumption and dividends follow a continuous, smooth process. However, the jumps calibrated in this paper are smaller and more frequent than in BCDG (they are 'infrequent' rather than 'rare'). Moreover, in this paper it is the desire for robustness against model uncertainty, calibrated using statistical detection probabilities, rather than the risk of being in the 'bad' regime, that induces a heightened fear of jumps and a demand for puts.

Shaliastovich (2011) also generates the implied volatility skew in a model with an Epstein-Zin representative agent and fundamentals driven by a long-run risks process. In his model the long-run risk component in cash flow growth is latent and the agent must learn it from a set of 'signals' received each period. The 'noisiness' of the signals (their cross-sectional standard deviation) is modeled as a persistent and volatile process that gets hit with jump shocks. In the model, periods in which the signals are noisier are characterized by volatile updates to the agent's filtered estimate. Hence, the noisiness of the signals acts like the volatility of the long-run risk component and the jump shocks to it cause jumps in prices, allowing the model to generate a high price for otm put options and a steep implied volatility skew. Shaliastovich (2011) does not, however, study the size of the embedded variance premium or its predictive power for stock returns. An issue with the approach followed by the model is that standard Bayesian learning implies that the filtered estimate for a latent variable will be *less* volatile than the variable itself, making it more difficult to generate high risk premiums in the model. More importantly, under Bayesian learning the conditional volatility of the latent state estimate varies inversely with signal noisiness, since high noise means the set of signals are less informative for updating. Therefore, with a Bayesian learner the state variable will update more smoothly during periods of high signal noise, implying *less* volatile stock returns and reducing the value of options. Instead, Shaliastovich (2011) uses a *non-Bayesian* representative agent who has 'recency bias' and sub-optimally puts excessive weight on recent observations.

Bollen and Whaley (2004) and Garleanu, Pedersen, and Poteshman (2009) argue that net buying pressure from investors for index options, coupled with an imperfect ability of market makers to hedge net option supply, lead to a high price premium for otm puts. This explanation takes the demand for options as exogenous. Hence, it does not explain why investors are willing to pay the high option premium or why, given the high (presumably risk-adjusted) returns available for selling options, other market participants, such as hedge funds, do not increase their option supply and thereby reduce the price premium.

A potentially interesting approach for option pricing that has not yet been implemented (to my knowledge) involves a representative agent with Disappointment Aversion preferences as in Gul (1991) or its extension to Generalized Disappointment Aversion (Routledge and Zin (2010)).[1] These preferences penalize outcomes where the agent's realized utility is below some fraction of his certainty equivalent. They have intuitive appeal for option pricing because put options can be used to help avoid the 'disappointing' outcomes. As shown in Routledge and Zin (2010), under Disappointment Aversion the agent's Euler equation involves a scaling up of the probabilities of disappointing states by a constant proportion. The drawback to this is that a constant probability scaling is not particularly well-suited to generating the implied volatility skew. To generate a steep skew, one would prefer a state-price density that continues to *increase* as outcomes become more negative, rather than staying constant. This is necessary to make puts that are further out-of-the-money increasingly valuable compared to Black-Scholes. This feature of the data is implied by the skew since the elasticity of put prices with respect to their Black-Scholes implied volatility *decreases* with their moneyness. Hence, the put price premium relative to Black-Scholes must increase as moneyness decreases, rather than stay flat, to generate a steep and monotonic skew. Generally speaking, since Disappointment Aversion is characterized by first-order risk-aversion, it "implies proportionately greater aversion to small risks than large ones" (Backus, Routledge, and Zin (2004)) and therefore does not seem particularly well suited for explaining otm put prices.

[1] I thank the Editor for suggesting a discussion of this approach.

II. Conditional Variance Forecast

I create a proxy for the conditional expectation of total return variation using the one-step ahead forecasts from a simple regression. I use the same regression specification used in Drechsler and Yaron (2009), extended to the longer sample. For the conditional forecast series, I use the one-step-ahead (e.g one-month ahead) forecast implied by the regression. The regression is estimated using the the full data sample, except for the rolling regression results in Table III. A brief summary is given below for completeness. See Drechsler and Yaron (2009) for a more detailed discussion.

Let Fut_t^2 and Ind_t^2 denote the realized variance in month t on the S&P 500 Futures and S&P 500 Index, respectively. The realized variance for a given month is found by summing up the squared five-minute log returns over the whole month. Five-minute log returns are calculated by taking the difference in the log price over 5 minute intervals. The conditional forecast is then the forecast of Fut_{t+1}^2 at time t. I follow Drechsler and Yaron (2009) by projecting it on lagged values of Ind^2 and VIX^2. The projection is estimated by OLS and the forecasts from it serve as the proxy for the series of conditional expectations of total return variation. The estimated coefficients are given below, with t-statistics in parenthesis, calculated using Newey-West (HAC) standard errors with 4 lags.

$$\begin{array}{cccccccccc}
\text{Fut}_{t+1}^2 & = & \alpha & + & \beta_1 & \text{Ind}_t^2 & + & \beta_2 & \text{VIX}_t^2 & + & \varepsilon_{t+1} & & \text{R}^2 \\
 & & 0.25 & & 0.35 & & & 0.54 & & & & & 0.46 \\
 & & (0.12) & & (3.82) & & & (5.99) & & & & &
\end{array}$$

The regression R^2 is 46%. Due to the large, rapid changes in realized variance during the financial crisis, the full-sample variance predictability is somewhat lower than during the pre-crisis period (about 59%). However, the coefficient estimates remain very similar to the ones estimated by Drechsler and Yaron (2009) in a pre-crisis period.

As explained in Secton II, I impose on the forecast proxy series the theoretical restriction that the physical expectation of total return variance be less than the risk-neutral expectation (which is given by VIX^2). I do this by truncating the value of the forecast series from above at the value of

VIX^2. This is equivalent to imposing nonnegativity on the estimated variance premium series. This is similar to Campbell and Thompson (2007), who argue in favor of imposing theoretical constraints on their forecasts. This restriction only has a noticeable impact on one of the 240 forecasts (the one for 2009.10).

Finally, for Table III, I use the one-step ahead forecasts from the same projection, but estimated on a rolling basis using only past data. Once again, I impose the same theoretical restriction using the same truncation scheme. The first 24 months of data are used to initialize the rolling regression estimates, so the proxy series begins in January 1992.

III. Derivation of Measure Changes

Recall that $\eta_t = \eta_t^{dZ}\eta_t^J$. I derive expressions for η_t^{dZ} and η_t^J corresponding to the alternative model dynamics discussed in the main text.

η_t^{dZ} solves the SDE $\frac{d\eta_t^{dZ}}{\eta_t^{dZ}} = h_t^T dZ_t$. An application of Ito's lemma shows that its solution is:

$$\eta_t^{dZ} = \exp\left(\int_0^t h_s^T dZ_s - \frac{1}{2}\int_0^t h_s^T h_s ds\right)$$

Note that η_t^{dZ} is a martingale and that $\eta_0^{dZ} = 1$. Let $P(\eta)$ be the measure that results from application of η to the reference measure P. Girsanov's theorem then implies that $Z_t^\eta = Z_t - \int h_t dt$ is a Brownian motion under $P(\eta)$. Writing the dynamics (1) in terms of Z_t^η alters the drift by adding to it the term $\Sigma(Y_t)h_t$, as in (2). This accounts for the perturbation to the drift under $P(\eta)$.

Since the Poisson process arrivals are (conditionally) independent and the jump sizes are i.i.d, the expression for η_t^J can be written as $\eta_t^{J_1}\eta_t^{J_2}\ldots$ where $\eta_t^{J_i}$ changes the probability law for the i-th jump component. I construct such terms to change the distribution of gamma-distributed jumps and normally distributed jumps.

Consider first gamma-distributed jumps, $\xi_i \sim \Gamma(k,\theta)$, where k and θ are the shape and scale parameters respectively. I want to construct the corresponding term $\eta_t^{J_i}$ in the Radon-Nikodym derivative so that under $P(\eta)$ the jump distribution is given by $\xi_i^\eta \sim \Gamma\left(k, \frac{\theta}{1-\theta b}\right)$, where b is the parameter that changes the gamma distribution's scale. I further specify the measure change so that the corresponding jump intensity changes from $l_{t,i}$ to $l_{t,i}^\eta = \exp(a) l_{t,i}$, i.e. it is scaled by the term $\exp(a)$ where a is a perturbation parameter. The desired $\eta_t^{J_i}$ solves the following SDE:

$$d\eta_t^{J_i} = \left(\exp\left[a + b\xi_i - \ln \psi_i(b)\right] - 1\right) \eta_{t\_}^{J_i} dN_t - \left(\exp(a) - 1\right) l_{t,i} \eta_t^{J_i} dt$$

where $\psi_i(b)$ is the moment-generating function of ξ_i evaluated at b. An application of Ito's lemma shows that $\eta_t^{J_i}$ is given by:

$$\eta_t^{J_i} = \exp\left(\int_0^t \left(a + b\xi_{i,s} - \ln \psi_i(b)\right) dN_s - \int_0^t l_{s,i} \left(\exp(a) - 1\right) ds\right)$$

Note that the process $\eta_t^{J_i}$ is a martingale and $\eta_0^{J_i} = 1$. Girsanov's theorem for jump processes then implies that under $P(\eta)$, the jump intensity is scaled by $\exp(a)$, as desired. Furthermore, under $P(\eta)$ the moment-generating function of ξ_i is given by:

$$\psi_i^\eta(u) = \frac{\psi_i^\eta(b+u)}{\psi_i(b)}$$

Straightforward substitution of the mgf for a gamma distribution shows that $\psi_i^\eta(u)$ is the mgf of a $\Gamma\left(k, \frac{\theta}{1-\theta b}\right)$, as desired.

Finally, I consider normally-distributed jumps, $\xi_k \sim \mathcal{N}(\mu, \sigma^2)$. I want to construct the corresponding term $\eta_t^{J_k}$ in the Radon-Nikodym derivative so that under $P(\eta)$ the jump distribution is given by $\xi_k^\eta \sim \mathcal{N}\left(\mu + \Delta\mu, \sigma^2 s_\sigma\right)$, where $\Delta\mu$ shifts the mean and s_σ scales the variance of the distribution. I further specify the measure change so that the corresponding jump intensity changes to $l_{t,k}^\eta = \exp(a) l_{t,k}$. The desired $\eta_t^{J_k}$ is given by:

$$\eta_t^{J_k} = \exp\left(\int_0^t \left(a + b_2\xi_k^2 + b_1\xi_k - \frac{1}{2}\left[\frac{(\mu+\Delta\mu)^2}{s_\sigma \sigma^2} - \frac{\mu}{\sigma^2} + \ln s_\sigma\right]\right) dN_s - \int_0^t l_{t,k} \left(\exp(a) - 1\right) ds\right)$$

where $b_1 = \frac{\mu(1-s_\sigma)+\Delta\mu}{s_\sigma \sigma^2}$ and $b_2 = \frac{1}{2}\frac{1}{\sigma^2}\left(1 - \frac{1}{s_\sigma}\right)$. By construction, the process $\eta_t^{J_k}$ is a martingale and $\eta_t^{J_k} = 1$.

Finally, since the the terms composing η_t are all martingales and have zero cross-variation (they are conditionally independent) η_t is a martingale with $\eta_0 = 1$ and therefore the measure $P(\eta)$ is indeed a probability measure.

IV. Equilibrium Consumption-Wealth Ratio

In equilibrium, markets clear so that the representative investor must hold all of his wealth in the aggregate consumption claim. To derive the equilibrium consumption-wealth ratio, consider the consumption and portfolio problem of the representative investor in this endowment setting. Under the reference measure, the price of the aggregate consumption claim P_c follows an Itô process of the form:

$$dP_{c,t} = (P_{c,t}u_{c,t} - C_t)dt + P_{c,t}\sigma_{c,t}^T dZ_t + P_{c,t-}(\exp(\Delta \ln P_{c,t}) - 1)$$

There is also a risk-free money market account in zero-net supply, paying an endogenously determined rate $r_{f,t}$. The investor chooses the proportion α_t of his wealth, W_t, to invest in the consumption claim. His budget constraint is then:

$$dW_t = W_t \left[\alpha_t(u_{c,t} - r_{f,t}) + r_{f,t}\right] dt + \alpha_t W_t \sigma_{c,t}^T dZ_t + \alpha_t(\exp(\Delta \ln P_{c,t}) - 1) - C_t dt$$

The lifetime utility of the investor $J(W_t, \tilde{Y}_t)$ is a function of W_t and the state variables for the dynamics, $\tilde{Y}_t$. The investor's HJB equation is:

$$0 = \max_{\{\alpha_t, C_t\}} \min_{P(\eta_t)} f(C_t, J_t) + E_t^{\eta}[dJ]$$

subject to the restriction on $R(\eta_t)$. We are interested in the investor's first-order condition with respect to C_t. Writing out the Lagrangian and taking the derivative with respect to C_t, the FOC is:

$$f_C(C, J) = J_W$$

Homogeneity of the preferences in wealth and linearity of the budget constraint imply that the value function must take the form $J(W, \tilde{Y}) = H(\tilde{Y})\frac{W^\gamma}{\gamma}$ for some function H. Substituting in for $f(C, J)$ and J_W their functional forms, simplifying, and rearranging, one obtains:

$$\frac{C}{W} = H(\tilde{Y})^{\frac{1-\psi}{\gamma}} \delta^\psi \tag{IA.1}$$

We want to obtain the consumption-wealth ratio in terms of the function $g(\tilde{Y})$. In equilibrium, the market clears and the investor consumes exactly the aggregate consumption stream, so lifetime utility is given by the equilibrium value of J in (6). Equating the two expressions for J and dividing through by W^γ gives:

$$H(\tilde{Y}) = \exp\left(\gamma g(\tilde{Y})\right) \left(\frac{C}{W}\right)^\gamma$$

Substituting this in for $H(\tilde{Y})$ in (IA.1) and solving for $\frac{C}{W}$ gives the result:

$$\frac{C_t}{W_t} = \exp\left(-\rho g(\tilde{Y}_t)\right) \delta \tag{IA.2}$$

V. Equity Return

I follow the approach of Eraker and Shaliastovich (2008). Let $\ln V_{t+s} = \ln M_{t+s} - \ln M_t + \int_t^{t+s} d\ln R_{m,u}$. The Euler equation implies that V_t is a martingale under the worst-case measure:

$$E_t^\eta[d\ln V_t^c + \frac{1}{2}(d\ln V_t)^2 + \exp(\Delta \ln V_t) - 1] = 0 \tag{IA.3}$$

where $d\ln V_t = d\ln M_t + d\ln R_{m,t}$. Log-linearizing $d\ln R_{m,t}$ around the unconditional mean of $v_{m,t}$ gives: $d\ln R_{m,t} = \kappa_{0,m}dt + \kappa_{1,m}dv_{m,t} - (1-\kappa_{1,m})v_{m,t}dt + d\ln D_t$. Further substituting into this expression the conjecture (11) for $v_{m,t}$ gives (12), which expresses $d\ln R_{m,t}$ in terms of $A_{0,m}$, A_m and the state variables. Substituting the expression for $d\ln R_{m,t}$ into $d\ln V_t$ along with the expression for $d\ln M_t$ (9) gives:

$$d\ln V_t = -\theta\delta dt - (1-\theta)\delta\exp\left(-\rho A_0 - \rho A'\tilde{Y}_t\right)dt + \kappa_{0,m}dt - (1-\kappa_{1,m})(A_{0,m} + A_m' Y_t)dt + \chi_m' dY_t$$

where $\chi_m = (-\Lambda + \kappa_{1,m}A_m + \delta_d)$. I now employ the exact same log-linearization of $\exp(-\rho A_0 + -\rho A'\tilde{Y}_t)$ given in Appendix Appendix C to replace it with $\kappa_0 + \kappa_1\rho A_0 + \kappa_1\rho A'\tilde{Y}_t$. Then substituting $d\ln V_t$ into (IA.3) and evaluating the expectation results in the following equation:

$$\begin{aligned} 0 = &-\theta\delta dt - (1-\theta)\left[\delta\kappa_0 + \delta\kappa_1\rho A_0\right]dt + \kappa_{0,m}dt - (1-\kappa_{1,m})A_{0,m}dt + \chi_m' E_t^\eta[dY_t^c] \\ &+ \left[(\theta-1)\delta\kappa_1\rho(\hat{A} - \delta_c) + (\kappa_{1,m}-1)A_m\right]' Y_t dt + \frac{1}{2}\chi_m^T\Sigma_t\Sigma_t^T\chi_m + l_t^{\eta\prime}\left(\psi^\eta(\chi_m) - 1\right) \end{aligned} \tag{IA.4}$$

We can now use the method of undetermined coefficients. This equation must hold for any value of Y_t, which implies that for each component in Y_t the sum of the terms multiplying it must be 0. Furthermore, the sum of the constant terms must be 0. Thus, the equation implies a system of $n+1$ equations whose solution is the $n \times 1$ vector A_m and the scalar $A_{0,m}$. The solution can be found numerically and verifies the conjectured functional form (11) for $v_{m,t}$.

VI. Integrated Variance and Risk-Neutral Dynamics

For convenience, let $* \in \{P, \eta, Q\}$ refer to either the reference, worst-case, or risk-neutral measure, respectively. Equation (12) implies that $E_t^*(d\ln R_{m,t})^2 = B_r'\Sigma_t\Sigma_t' B_r + B_r^{2\prime}\left[E^*(\xi_t^2)\cdot l_t^*\right]$, where B_r^2 denotes the vector obtained by squaring the components of B_r. We want to calculate the expectation of integrated variance: $E_t^*[\int_t^T (d\ln R_{m,s})^2] = \int_t^T E_t^*\,(d\ln R_{m,s})^2$. To that end, it is useful to write $E_t^*(d\ln R_{m,t})^2 = \alpha_0^* + \alpha^{*\prime}Y_t$ where α_0 is a scalar and α is a vector of loadings on the state Y_t. The law of iterated expectations implies that: $\int_t^T E_t^*\,(d\ln R_{m,s})^2 = \int_t^T(\alpha_0^* + \alpha^{*T}E_t^*(Y_t))$

A straightforward expansion of the expression for $E_t^*(d\ln R_{m,t})^2$ shows that:

$$\begin{aligned}\alpha_0^* &= B_r' h B_r \\ \alpha^{*\prime} &= B_r' H B_r + B_r^{2\prime}\mathrm{diag}(E^*(\xi_t^2)) l_1^*\end{aligned}$$

where $B_r' H B_r$ denotes a row vector where the i-th component is $B_r' H_i B_r$. From this we see that only α^* differs across the measures.

In order to calculate expectations of future values of Y_t, which is required to calculate the integral, it is easiest to express the dynamics of Y_t in terms of demeaned jump shocks (i.e. using the 'compensated' Poisson processes). The general form of compensated dynamics is:

$$dY_t = \mu^* + \widehat{\mathcal{K}^*} Y_t + \Sigma_t dZ_t^* + \xi_t^* \cdot dN_t^* - E_t^*(\xi_t^* \cdot dN_t^*)$$

where $\widehat{\mathcal{K}^*}$ is the resulting transition matrix which incorporates the uncompensated transition matrix, $\mathcal{K}^*$, and the compensation to the jump terms. A standard calculation then gives that:

$$E_t^*(Y_{t+\Delta t}) = \exp(\widehat{\mathcal{K}^*}\Delta t) Y_t + \widehat{\mathcal{K}^*}^{-1}\left(\exp(\widehat{\mathcal{K}}^*\Delta t) - I\right)\mu^*$$

where I is the identity matrix. A straightforward calculation of the integral in $\int_t^T(\alpha_0^* + \alpha^{*T} E_t^*(Y_t))$ results in the following expression for expected integrated variance:

$$E_t^*\left[\int_t^{t+\Delta t}(d\ln R_{m,s})^2\right] = \alpha_0^*\Delta t + \alpha^{*\prime}\left[\Theta Y_t + \widehat{\mathcal{K}^*}^{-1}\left[\Theta - I\cdot\Delta t\right]\mu^*\right] \qquad \text{(IA.5)}$$

where $\Theta = \widehat{\mathcal{K}^*}^{-1}(\exp(\widehat{\mathcal{K}}^*\Delta t) - I)$.

Finally, I derive the parameters of the compensated dynamics under the three measures. Recall that $\psi^*(u)$ denotes the stacked vector of moment-generating functions evaluated at the vector u. Then we have that $E_t^*(\xi_t^* \cdot dN_t^*) = \mathrm{diag}(\psi^{*(1)}(0)) l_1^* q_t^2$ where $\psi^{*(1)}(0)$ is the first derivative of $\psi^*(u)$ evaluated at 0. Let δ_q be the selector vector for q_t^2, i.e. $\delta_q' Y_t = q_t^2$. Then denote by

$[l_1^*]_q$ a matrix that has the vector l_1^* in the 'q-th' column, so that $[l_1^*]_q \delta_q = l_1^*$, and all other columns set equal to the 0 vector. The transition matrix for the compensated dynamics under P is then given by: $\widehat{\mathcal{K}^P} = \mathcal{K} + \mathrm{diag}(\psi^{(1)}(0))[l_1]_q$ and $\mu^P = \mu$. Under the worst-case model: $\widehat{\mathcal{K}^\eta} = \mathcal{K} + [\Sigma_t h_t / q_t^2]_q + \mathrm{diag}(\psi^{\eta(1)}(0))[l_1^\eta]_q$ and $\mu^\eta = \mu$. The difference from P comes from the drift perturbation and the change in the jump intensity and moment-generating function. Finally, under the risk-neutral measure, $\mathcal{K}^Q = \mathcal{K} + [\Sigma_t h_t / q_t^2]_q - H\Lambda + \mathrm{diag}(\psi^{Q(1)}(0))[l_1^Q]_q$ and $\mu^Q = \mu - h\Lambda$, where $H\Lambda$ denotes an $n \times n$ matrix with k-th column equal to $H_k\Lambda$.

The risk-neutral moment-generating function and jump intensity are determined by the worst-case moment-generating function and jump intensity and the price of risk vector Λ. The moment-generating functions are given by $\psi^Q(u) = \psi^\eta(-\Lambda + u)/\psi^\eta(-\Lambda)$, where the division is component-wise. The jump intensity vector is $l_q^Q = \psi^\eta(-\Lambda) \cdot l_t^\eta$. For a proof see Proposition 5 in Duffie, Pan, and Singleton (2000). The risk-neutral expressions show that in going from η to Q, the change in jump intensities and distributions depends on the prices of risk Λ. Risk-neutralization tilts probabilities towards 'high-price' states of the world. The direction and amount of the 'tilt' depends on the magnitude of Λ. For example, note that if $\Lambda = 0$ the worst-case and risk-neutral quantities are identical.

VII. Detection Error Probabilities

Detection error probabilities are a useful tool for calibrating model uncertainty that is due to Anderson, Hansen, and Sargent (2003). The detection error probability gives the probability that, using a likelihood-ratio test, a decision maker will incorrectly reject the worst-case model in favor of the reference model based on a data sample of a given length T. This is an important statistic because the investor is exactly worried about the possibility that the data has led him to favor the reference model although the true data-generating process is the worst-case model. I now explain how the detection error probability can be calculated in terms of the Radon-Nikodym process η_t.

The likelihood ratio of the worst-case model to the reference model is exactly given by the

Radon-Nikodym derivative η_t. Therefore, the probability at time zero of making a detection error based on a sample of length T is $\text{Prob}^{\eta}(\ln \eta_T < 0|\mathcal{F}_0, \eta_0 = 1)$. Note that the probability is evaluated under the worst-case measure. For illustration, I derive the detection error probability for an i.i.d pure diffusion reference model and then discuss how it can be calculated for the framework in this paper.

As Appendix III shows, in a pure diffusion setting $\ln \eta_T = \int_0^T h_t^T dZ_t - \frac{1}{2}\int_0^T h_t^T h_t dt$. Substituting in $dZ_t = dZ_t^{\eta} + h_t dt$ gives an expression that is more convenient for evaluation under the worst-case measure:

$$\ln \eta_T = \int_0^T h_t^T dZ_t^{\eta} + \frac{1}{2}\int_0^T h_t^T h_t dt$$

Now consider the distribution of η_T under the worst-case measure. Taking expectations gives

$$E_0^{\eta}[\ln \eta_T] = \frac{1}{2}\int_0^T E_0^{\eta}\left[h_t^T h_t\right] dt = \frac{1}{2}\int_0^T 2\varphi = \varphi T$$

When the reference and worst-case models are i.i.d, h_t is constant. It then follows that $\ln \eta_T$ has a normal distribution with variance $T \times h^T h = 2\varphi T$, i.e. $\ln \eta_T \overset{\eta}{\sim} \mathcal{N}(\varphi T, 2\varphi T)$. The detection error probability is then:

$$\text{Prob}^{\eta}(\ln \eta_T < 0|\mathcal{F}_0, \eta_0 = 1) = \text{Prob}\left(\mathcal{N}(0,1) < \frac{-\varphi T}{\sqrt{2\varphi T}}\right) = \text{Prob}\left(\mathcal{N}(0,1) < \frac{-1}{\sqrt{2}}\sqrt{\varphi T}\right)$$

Therefore, in this simple case, the detection error probability is $\Phi(\frac{-1}{\sqrt{2}}\sqrt{\varphi T})$, where Φ is the cdf of the standard normal distribution.

In general, a closed-form expression for the detection error probability is not available since the distribution of $\ln \eta_T$ is not known in closed-form. However, for a general class of specifications that includes the affine setting of this paper, the detection error probability can be calculated numerically via Fourier inversion. As Maenhout (2006) shows, using the expression for η_T from Appendix III, one can find the (conditional) characteristic function of η_T in closed-form (up to a system of ODEs). The exact detection error probability can then be calculated numerically via a Fourier inversion. This methodology is similar to the one used to calculate option prices in affine settings, as developed in Duffie, Pan, and Singleton (2000). Maenhout (2006) contains a detailed derivation. In calculating the detection error probabilities for the calibrated model, I set the time-0 value of the state vector equal to its unconditional mean. Except for short samples, the detection-probabilities are relatively insensitive to the time-0 value of the state vector.

VIII. An Option-Extracted Uncertainty Series

I obtain a time series of q_t^2 and σ_t^2 from the empirical options data by using the model-based option prices. I do this by finding the values of these two state variables that provide the best fit to a cross-section of implied-volatilities at each date t using the calibrated model. Although technically x_t also effects option prices via the risk-free rate, this effect is minuscule and should be swamped by any noise in the price data. I therefore fix x_t at its unconditional mean. Using at least two implied-volatilities on a given date t, one can solve for the values of the two state variables that generates the best model fit to the data. While using exactly two implied-volatilities generates an exact fit, one can include additional options and do a (non-linear) least-squares fit in order to attenuate the impact of any data noise.

Figure IA.3 shows the extracted series obtained by fitting the model-based implied volatilities to their empirical counterparts for strikes with moneyness of 1, 0.9, and 0.8 at a 1-month maturity. For each month, I use the implied volatilities for the last day of the month. The series were obtained from CSFB and span 1998.9 to 2008.7. A moneyness of 0.8 represents put options that are far out-of-the-money and should be informative regarding investor's fears of negative jump shocks.

The figure shows that the implied q_t^2 series is volatile and is occasionally hit by extreme spikes, while the implied σ_t^2 series is smoother. The extreme spikes in q_t^2 correspond to the the periods of the 1998 LTCM crisis, September 11th, and the corporate scandals of 2002. Both series are low and tranquil for a period that begins in 2004 and ends in 2007. The time-series properties of the series appear to be overall consistent with those implied by the calibrated model. The means of the implied $(\sigma_t^2,\, q_t^2)$ are $(1.17, 1.04)$, which is close to their population means of 1, while their standard deviations are $(0.64, 1.12)$, which is close to the population values of $(0.69, 1.19)$. The autocorrelations of the implied (σ_t^2, q_t^2) are 0.80 and 0.68 respectively, which are somewhat lower than their population counterparts, while the correlation between the two series is 0.15.

Figure IA.1: Implied Volatilities: Model and Data

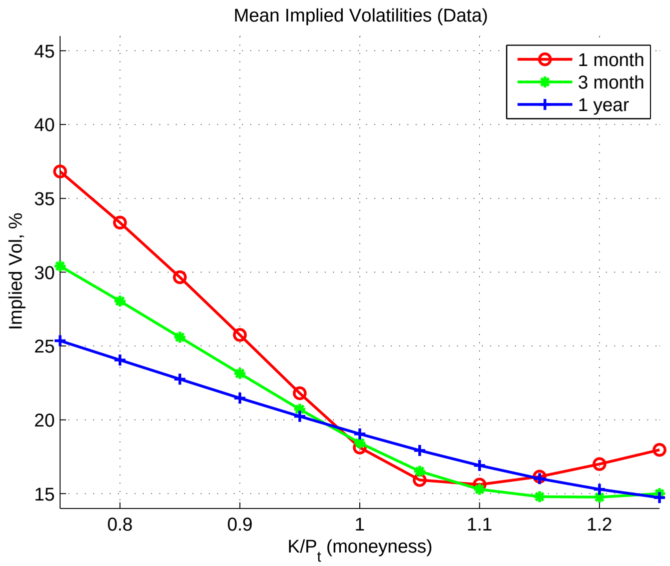

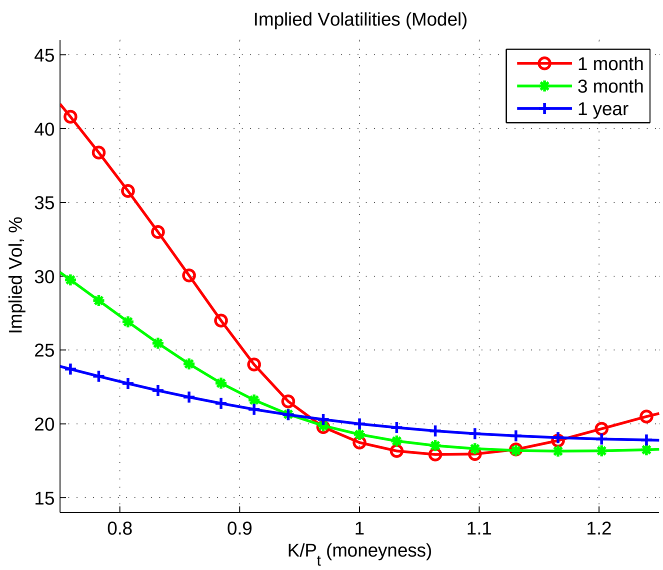

The figure plots implied volatilities from empirical option prices and for option prices calculated for the model of Table IV. The plots show implied-volatility curves for maturities of 1, 3, and 12 months. Strikes are expressed in moneyness (Strike Price/Spot price). The top plot shows the mean of daily implied volatilities for S&P 500 index options for the period 1999.10-2008.6, as quoted in the over-the-counter market. The bottom plot shows the model-based implied volatilities for option prices obtained when the model's state vector is set equal to its unconditional mean.

Figure IA.2: Implied Volatilities: No-Uncertainty Model and Data

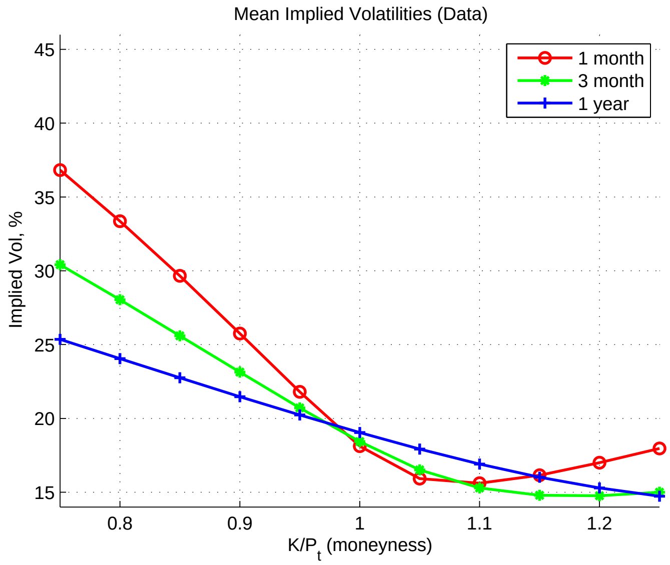

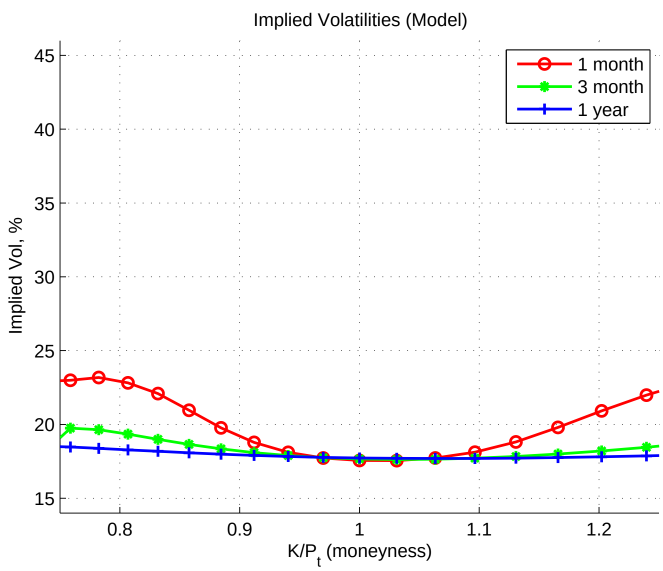

The figure plots the implied volatilities from empirical option prices and for option prices calculated for Model 1-B, which was used in the comparative statics exercise in Table VIII. The plot shows curves for maturities of 1, 3, and 12 months. Strikes are expressed in moneyness (Strike Price/Spot price). The top plot shows the mean of daily implied volatilities for S&P 500 index options for the period 1999.10-2008.6, quoted in the over-the-counter market. The bottom plot shows the model-based implied volatilities for option prices obtained when the model's state vector is set equal to its unconditional mean.

Figure IA.3: Option-Implied Time Series of State Variables

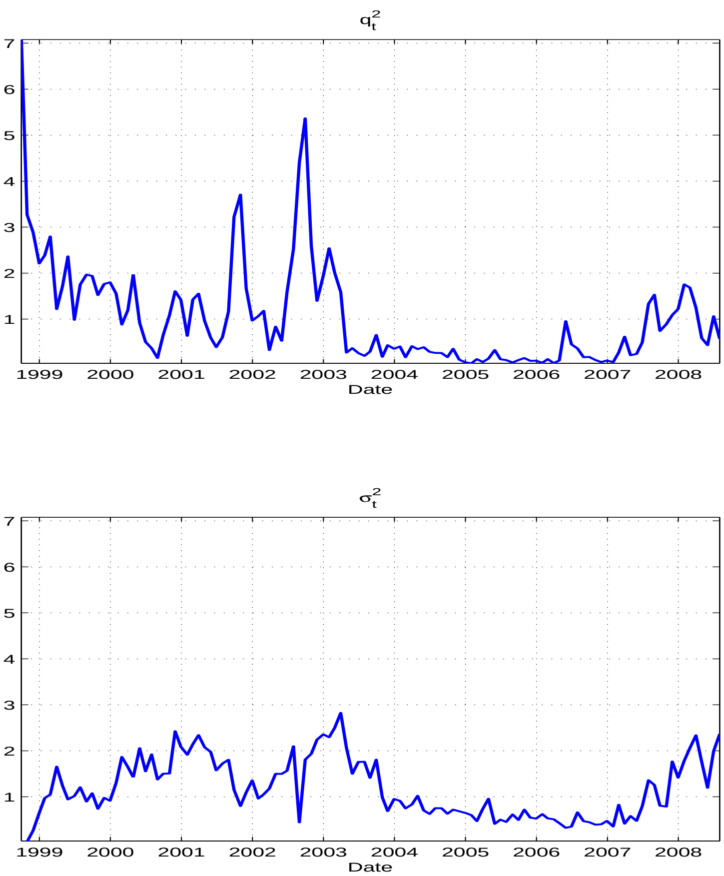

The figure plots the time-series of q_t^2 and σ_t^2 extracted from empirical option prices using the model of Table IV. The sample is 1998.9-2008.7.

References

Anderson, Evan, Lars P. Hansen, and Thomas J. Sargent, 2003, A Quartet of Semigroups for Model Specification, Robustness, Prices of risk, and Model Detection, *Journal of the European Economic Association* 1, 68–123.

Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters Implied by Equity Index Options, *Journal of Finance* forthcoming.

Backus, David, Bryan Routledge, and Stanley Zin, 2004, Exotic Preferences for Macroeconomists, *NBER Macroeconomics Annual*.

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

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.

Bollen, Nicolas, and Robert Whaley, 2004, Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?, *Journal of Finance* 59, 711–753.

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, 205–255.

Campbell, John Y., and Samuel B. Thompson, 2007, Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?, *Review of Financial Studies* forthcoming.

Drechsler, Itamar, and Amir Yaron, 2009, What's Vol Got To Do With It, Working paper, The Wharton School, University of Pennsylvania.

Du, Du, 2010, General Equilibrium Pricing of Options With Habit Formation and Event Risks, *Journal of Financial Economics* 99.

Duffie, D., J. Pan, and K. J. Singleton, 2000, Transform Analysis and Asset Pricing for Affine Jump-Diffusions, *Econometrica* 68, 1343–1376.

Eraker, Bjorn, and Ivan Shaliastovich, 2008, An Equilibrium Guide to Designing Affine Pricing Models, *Mathematical Finance* forthcoming.

Gabaix, Xavier, 2011, Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance, Working paper, New York University.

Garleanu, Nicolae, Lasse H. Pedersen, and Allen Poteshman, 2009, Demand-Based Option Pricing, *Review of Financial Studies* 22, 4259–4299.

Gul, Faruk, 1991, A theory of disappointment aversion, *Econometrica* 59, 667–686.

Maenhout, Pascal J., 2006, Robust Portfolio Rules and Detection Error Probabilities for a Mean-reverting Risk Premium, *Journal of Economic Theory* 128, 136–163.

Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding predictability, *Journal of Political Economy* 112(1), 1–47.

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

Routledge, Bryan R., and Stanley E. Zin, 2010, Generalized Disappointment Aversion and Asset Prices, *Journal of Finance* 65, 1303–1332.

Shaliastovich, Ivan, 2011, Learning, Confidence, and Option Prices, Working paper, Wharton School, University of Pennsylvania.

Wachter, Jessica, 2011, Can time-varying risk of rare disasters explain aggregate stock market volatility?, Working paper, University of Pennsylvania.