Inventory Risk, Market-Maker Wealth, and the Variance Risk Premium: Theory and Evidence

March 13, 2015

Abstract

We investigate the role of option market-makers in determining the variance risk premium. A substantial part of the variance risk premium is driven by market-makers' risk sharing capacity. When market-makers experience dramatic wealth losses, a one-standard-deviation change in inventory risk leads to more than 9% change in the variance risk premium. Motivated by our findings, we develop a model in which a market-maker with limited capital is exposed to market variance risk through inventory. We derive an endogenous variance risk premium and analyze its dependence on inventory risk and market-maker wealth. Estimating the model on index returns and options we find that it fits the data well specifically during the financial crisis.

JEL Classification: G10; G12; G13.

Keywords: Variance risk premium; inventory risk; financial constraints; option pricing.

1 Introduction

The variance of the stock market evolves stochastically over time. Consequently, it exposes market participants to market variance risk in addition to market return risk. Market variance is thus another source of aggregate risk and commands a premium for its exposure. This premium impacts the return and price of any security that co-moves with market variance. It influences the cross-section of stock expected returns, it predicts the returns on the stock market, and it impacts the cross-section of index option prices (see, among others, Bakshi, and Kapadia, 2003, Ang, Hodrick, Xing, and Zhang, 2006, and Bollerslev, Tauchen, and Zhou, 2010). Because the variance risk premium has important pricing implications, a better understanding of the factors influencing its determinants is of significant interest.

With more than one hundred million dollars of daily volume, index options are the most actively traded securities that provide direct exposure to market variance risk. This makes the trading activity in index options particularly informative about the variance risk premium. In this study, we model index option market-makers and find, empirically and theoretically, a substantial impact of fluctuations in intermediaries' risk bearing capacity on the variance risk premium.

On average, the net demand by end-users of index options is positive (e.g., Bollen and Whaley, 2004, and Gârleanu, Pedersen, and Poteshman, 2009). Market-makers in these options act as net sellers and build up large negative inventories over time. Through their inventory, intermediaries are exposed to market variance risk in addition to market return risk. While market-makers efficiently hedge return risk using index future contracts, frictions impair their ability to eliminate their large exposures to market variance (see Bates, 2003). Consequently, market-makers bear large variance risk and their risk sharing capacity influences the premium for market variance.

We identify two variables informative about market-makers' capacity to intermediate variance risk that help explain the variance risk premium. The first variable captures the exposure of option market-makers' inventory to market variance risk. We refer to it as inventory risk. The second variable corresponds to market-makers' wealth, which we measure using aggregate trading revenues.

Gârleanu, Pedersen, and Poteshman (2009) show that net option demand exerts pressure on option prices when markets are incomplete. Their study sets the stage for our analysis of the way inventory risk and market-makers' wealth interact to impact the variance risk premium. We model a continuous-time economy in which the market variance has its own dynamic, and a risk-averse representative market-maker with limited capital provides quotes for index option. Because market variance fluctuates randomly over time, the market-maker is exposed to variance risk which we assume unhedgeable. We solve for the endogenous variance risk premium that induces the intermediary to clear the index option market. The solution derived provides new insights.

First, inventory risk and the variance risk premium co-move to reward the market-maker for his risk exposure. When the market-maker steps-in to absorb net buying orders, his inventory becomes more negative and so does his inventory risk. Because the market-maker is risk-averse, he requires a higher compensation which, given his negative exposure, translates into a more negative variance risk premium. Second, the model uncovers the substantial effect the market-maker's wealth has on the variance risk premium. When the intermediary incurs losses, his marginal utility of wealth rises increasing his required compensation. When the market-maker bears negative variance exposure a higher compensation implies a more negative variance risk premium. Thus, our study complements Gârleanu, Pedersen, and Poteshman (2009) by deriving an explicit relation between market-maker wealth, inventory risk and the variance risk premium in a stochastic volatility model.

We test the model's equilibrium predictions using fifteen years of market-making activity for index options and more than one million quotes. Both inventory risk and changes in market-maker wealth significantly explain the variance risk premium. When inventory risk becomes more negative by one standard deviation, it results in 1.2% decrease in the variance risk premium which is more than twenty times the average daily change in the variance risk premium. Moreover, this effect is magnified when market-makers experience dramatic wealth losses. When market-makers' loss is at its ninetieth percentile, a one standard deviation decrease in inventory risk can cause up to 9% decrease in the variance risk premium.

To further assess the robustness of our results, we investigate the ability of inventory risk and changes in market-maker wealth to explain daily changes in the variance risk premium measured at various horizons. While robust to the entire term structure of premia, the impact of inventory risk and market-maker wealth is most prominent at short-horizons.

Bollen and Whaley (2004) demonstrate that net buying pressures positively impact option implied volatilities. However, the channels by which net demand impacts option prices, and thus implied volatilities, remain unclear. Our analysis complements their study by identifying the intermediary channel. Net buying pressures lead to a more negative inventory risk which results in a more negative variance risk premium. Ultimately this decrease in variance risk premium translates into a higher risk-neutral variance and more expensive index options.

This paper is related to a growing literature on the variance risk premium. Early studies investigating the properties of this premium include Bakshi and Kapadia (2003), Vilkov (2008) and Carr and Wu (2009).[1] More recently, Egloff, Leippold, and Wu (2010) and Todorov (2010)

[1]See also Bakshi and Kapadia (2006) who relate the variance risk premium to investor risk aversion and the higher-order moments of the return distribution. Moreover, Driessen, Maenhout, and Vilkov (2009) explain the variance risk premium by the correlation risk of individual equity.

explain the variance risk premium dynamic using stochastic volatility models with jumps. Ang, Hodrick, Xing, and Zhang (2006) and Cremers, Halling, and Weinbaum (2014) show that the price of variance risk helps explain the cross-section of expected stock returns. Barras and Malkhozov (2014) study the difference between the variance risk premium estimated from the cross-section of stock returns and the one implied by options. We instead model the role of option intermediaries in the determination of the variance risk premium.

The recent financial crisis underlines the non-trivial influence of financial intermediaries' positions and constraints on asset prices (see, for instance, Adrian and Shin, 2010). In standard stochastic volatility models, the variance risk premium is the product of a time-invariant parameter and the latent spot variance (see, among others, Heston, 1993, Bates, 2000, Pan, 2002, and Egloff, Leippold, and Wu, 2010). This specification attributes any discrepancy between the objective distribution of index returns and the risk-neutral probability measure implied by option prices to the marginal investor's preference over market variance. Our study departs from this strand of literature by modeling the influence of market-makers' risk bearing capacity on the variance risk premium. This feature is related to Adrian, Etula, and Muir (2014) who show that intermediaries' leverage ratio is important for explaining the cross-section of equity returns.

One key strength of the model we develop is its parsimony. We exploit this attribute and estimate the model on index returns and a large panel of index put options. Relative to Heston (1993), the model performs well both in- and out-of-sample. More interestingly, it performs particularly well during the crisis and when pricing out-of-the-money puts, which tend to challenge traditional models. During turbulent times, our estimation suggests that fluctuations in market-makers' wealth can lead to more than 2% daily change in option prices. Under perfect integration of financial markets, intermediaries' risk exposure and wealth should not influence equilibrium prices. The evidence documented in this study implies a certain segmentation of the option market consistent with a limits to arbitrage argument. We refer to Shleifer and Vishny (1997), Gromb and Vayanos (2002), and Brunnermeier and Pedersen (2009) among others for prominent examples of such theories.

Related to our study is the work of Leippold and Su (2011) who examine the influence of option contracts' margin requirements on option implied-volatilities in a constant volatility framework. Also linked to our work is the study of Chen, Joslin, and Ni (2013). They investigate the jump premium embedded in index options and its predictive ability for stock market returns. In their model, consumption has a constant diffusive volatility and time-varying jumps. Moreover, intermediary constraints are modeled exogenously through time-varying risk-aversion. Our study complements both of these papers by designing a model in which market variance is stochastic, and market-maker wealth is endogenous. Both features are needed to obtain the explicit implications of the impact of intermediaries' risk bearing capacity on the premium for market variance.

The remainder of the paper is organized as follows. Section 2 presents the data and the main variables used in this study. It also presents our main empirical tests. In Section 3, we develop the theoretical model. Section 4 discusses the model implications. In Section 5, we present the model estimation. Section 6 concludes.

2 Empirical Analysis

2.1 Data

In our empirical analysis, we focus on the most liquid contracts providing direct exposure to market variance, S&P 500 index options (SPX options). SPX options trade only on the Chicago Board Option Exchange (CBOE). To construct the aggregate inventory of CBOE market-makers, we rely on the Market Data Express Open/Close database for obtaining the daily non market-makers (end-users) order flows for SPX options starting on January 1, 1996 and ending on December 30, 2011. The data divides the daily open/close buy and sell origins into two groups: firm and customer. Customer origin is further subdivided in three categories depending on order size. Since we are interested in measuring market-makers' aggregate inventory, we disregard individual origins and group orders into buy and sell. On each day we aggregate the total buys and sells for each contract and take the difference which corresponds to end-users' net demand for that contract on that day. Because SPX options are European, the time-series of market-makers' inventory for each contract can be reconstructed by summing up the negative of end-users' net demand over time starting from the first day the contract is quoted. We filter out Leaps (options with more than one year to maturity) and restrict the sample to start on January 1, 1997 to avoid biases in inventory measurement.[2]

For SPX option prices, we rely on end-of-day data from OptionMetrics starting on January 1, 1997 and ending on December 30, 2011. From all the quotes provided, we filter out options that have moneyness (spot price over strike price) less than 0.8 and larger than 1.2, those with a mid-quote less than 3/8, those with implied volatility less than 5% and greater than 150%, and those with less than ten days to maturity. The price of each contract is defined by the bid-ask mid-quote. For each option maturity, interest rates are estimated by linear interpolation using zero coupon Treasury yields. The dividend yield is obtained from OptionMetrics. We then merge the inventory data with the OptionMetrics database. The final sample contains more than one million quotes for

[2] To correctly measure inventory for a given option, the full time-series of end-users' order flows for that option must be observed. Because some options quoted during 1996 have been issued during 1995, accurate measurement of inventories constrains the empirical analysis to start on January 1, 1997.

SPX puts and calls over the 1997-2011 period.

To construct the variance risk premium, we need daily estimates of realized variance. To this end, we obtain high frequency data for S&P 500 index futures from Tickdata starting on January 1, 1997 and ending on September 30, 2012.[3] Based on multiple grids of intraday squared returns, we construct daily measures of average realized variance following Zhang, Mykland, and Aït-Sahalia (2005).

2.2 Definition of Variables

Variance Risk Premium

The variance risk premium captures the difference between physical and risk-neutral market variances. At time t the (annualized) variance risk premium with T-days horizon is

$$VRP_{t,T} \equiv RV_{t,T} - RNV_{t,T}, \tag{2.1}$$

where $RV_{t,T} \equiv E_t^P[\frac{1}{T}\int_t^{t+T} V_s ds]$ denotes the expected integrated physical variance and $RNV_{t,T} \equiv E_t^Q[\frac{1}{T}\int_t^{t+T} V_s ds]$ is the expected integrated risk-neutral variance. Suppose that we want to obtain a model-free estimate of the one-month variance risk premium on day t, that is $VRP_{t,30}$. Our first step is to compute $RV_{t,30}$. As in Carr and Wu (2009) and Egloff, Leippold, and Wu (2010), we proxy expected physical variance by ex-post realized variances as[4]

$$RV_{t,30} = \frac{365}{30} \cdot \left(\sum_{i=1}^{30} RV_{t+i-1} \right). \tag{2.2}$$

To measure expected integrated risk-neutral variance, we follow Britten-Jones and Neuberger (2000) and Bollerslev, Tauchen, and Zhou (2009) among others. We compute $RNV_{t,30}$ from a portfolio of SPX call options as

$$RNV_{t,30} = \frac{365}{30} \cdot \left(2 \cdot \int_0^\infty \frac{C(t, 30, Ke^{-r \cdot 30/365}) - C(t, 0, K)}{K^2} dK \right), \tag{2.3}$$

where $C(t, T, K)$ is the price of a call observed at time t with T days to maturity and strike price K, and r is the risk-free rate. We evaluate (2.3) using the trapezoidal rule. Armed with the two measures of expected integrated variance, the one-month variance risk premium on day t is

[3] For some of our empirical tests, we need estimates of realized variance up to September 30, 2012 to construct a measure of the 9-month ex-post realized variance on December 30, 2011.

[4] The results documented in this study are robust to the use of a predictive model in order to measure future expected realized variance. We discuss this further later.

$VRP_{t,30} = RV_{t,30} - RNV_{t,30}$.

Table 1 presents the daily average of implied volatility, vega, days to maturity, number of quotes, and volume.[5] We also report the yearly averages of the one-month variance risk premium. Market implied volatility is 22.28% on average. As found in earlier studies (see, among others, Bakshi and Kapadia, 2003, Vilkov, 2008, and Carr and Wu, 2009), the yearly averages of the variance risk premium are consistently negative. In our sample, the risk-neutral variance exceeds the realized variance by 2.13%.

Figure 1 plots the time-series of the S&P 500 index level in the top panel, and of the one-month variance risk premium expressed in percentage in middle panel. As expected, the variance risk premium has greater variations during periods of high uncertainty relative to periods of low volatility. This explains the relatively small variance risk premium during the 2003-2007 sample period. Note the way the variance risk premium is mostly negative before 2008. Interestingly, during the financial crisis it spikes to large positive and negative values contemporaneously to the abrupt price drop in the S&P 500 index. To investigate if the large variations in variance risk premium during the financial crisis are induced by measurement errors, we plot the weekly averages of daily gains and losses of delta-hedged near-the-money options in the bottom panel of Figure 1. This exercise is motivated by Bakshi and Kapadia (2003), who show that the gains and losses from delta-hedged positions in options are informative about the variance risk premium. For a given option f_t^j, the daily dollar gains and losses from delta-hedging it from day $t-1$ to t is

$$\Delta Hedge_t^j \equiv f_t^j - \left(\Delta_{t-1}^j S_t + (f_{t-1}^j - \Delta_{t-1}^j S_{t-1}) \cdot (1 + r\Delta t) + \Delta_{t-1}^j S_{t-1} \cdot q\Delta t\right), \tag{2.4}$$

where $\Delta_t^j \equiv \frac{\partial f_t^j}{\partial S_t}$ is option j delta, S_t denotes the value of the S&P 500, q is the dividend yield, and the time-step $\Delta t = 1/365$. On each week, we average the daily $\Delta Hedge_t^j$ for all options for which $0.98 \leqslant S_t/K^j \leqslant 1.02$ to obtain the weekly average gain and loss. Interestingly, the large fluctuations in $VRP_{t,30}$ during the crisis period are also apparent from the time-series of delta-hedged gains and losses.

Inventory Risk and Market-Maker Wealth

About 273 SPX calls and puts with distinct moneyness and maturities are quoted on each day. To assess market-makers' inventory across options, Table 2 reports the daily average of implied volatility, inventory, and delta-hedged gains and losses for different moneyness and maturity categories. Note how market-makers' positions are consistently negative across moneyness and maturities. Intermediaries are short about one hundred thousand contracts on a daily basis.

[5] The vega of an option measures the rate of change of the option price to a small increase in volatility.

In our analysis, inventory risk captures the aggregate exposure of index option market-makers' inventory to market volatility. At any time t, it is defined by

$$InvRisk_t \equiv \sum_j \Phi_t^{MM,j} \cdot Vega_t^j, \tag{2.5}$$

where $\Phi_t^{MM,j}$ denotes market-makers' inventory for option j, and $Vega_t^j \equiv \frac{\partial f_t^j}{\partial \sqrt{V_t}}$ is the option vega. The option vega is the sensitivity of the option price to market volatility. Accordingly, the aggregate exposure of market-makers to market variance is the sum of their inventory times vega across all contracts. Inventory risk is highly informative about market-makers' exposure as $1\% \times InvRisk$ indicates the way inventory would respond in dollar terms to a 1% increase in market volatility. From its definition, we see that inventory risk is signed. Because vega is always positive, it is the commonality in inventory across contracts that determines the sign of inventory risk. At times when intermediaries act as net sellers and $\Phi_t^{MM,j} < 0$ for most j, we have $InvRisk_t < 0$. In contrast, inventory risk is positive when market-makers carry long positions on average. Empirically, we calculate inventory risk on each day using SPX market-makers' positions combined with Black-Scholes' vega.[6]

Figure 2 plots the VIX index in the top panel, and the dynamic of inventory risk in the bottom panel. Interestingly, no clear patterns are apparent between the VIX and inventory risk time-series. Consistent with Table 2, we see that market-makers' exposure to S&P 500 volatility is negative most of the time excepted during the financial crisis, and in the beginning of 2011. Through their inventory, option market-makers carry billion dollars in risk exposure to market variance.

In order to measure the changes in market-makers' wealth over time, we first compute the daily profits and losses derived by market-makers from carrying their hedged inventory

$$P\&L_t \equiv \sum_j \Phi_{t-1}^{MM,j} \cdot \Delta Hedge_t^j. \tag{2.6}$$

where $\Delta Hedge_t^j$ satisfies (2.4). At the close of each day, the market-makers' aggregate daily profits and losses is the sum of lagged inventory times the delta-hedge gains and losses realized on that day across all contracts. The existence of bid-ask spreads implies another source of revenue for market-makers. Accordingly, we estimate the daily bid-ask spread revenue earned by market-makers by

[6] Black-Scholes' vega has also been used as a proxy for option vega in a stochastic volatility framework in Carr and Wu (2007), and Trolle and Schwartz (2009) among others.

calculating for each day

$$BA_t \equiv \sum_j \min(BO_t^j, SO_t^j) \cdot \left(BidAsk_t^j - 0.36\right), \tag{2.7}$$

where BO_t^j denotes end-users' buy orders for option j, SO_t^j is end-users' sell orders, $BidAsk_t^j$ denotes the option bid-ask spread, and \$0.36 is the transaction fee charged to dealers per contract traded.[7] In our empirical analysis, we define changes in wealth as the sum of (2.6) and (2.7), $\Delta W_t = P\&L_t + BA_t$. Our definition of market-makers' wealth is related to the measure used in Comerton-Forde, Hendershott, Jones, Moulton, and Seasholes (2010) to proxy NYSE specialists' revenues. However, our measure differs from theirs to account for the common practice of inventory delta-hedging adopted by option intermediaries.

Figure 3 plots the daily profits and losses from market-makers' delta-hedged inventory in the top panel, the cumulative daily profits and losses in the middle panel, and the bid-ask spread revenue in the bottom panel. Market-makers face substantial risks. Their daily profits and losses fluctuate between +265 and −393 million dollars. Consistent with the idea that large positions in options involve risks, the distribution of the daily delta-hedged gains and losses is highly leptokurtic with an excess kurtosis of 73. Moreover, it is asymmetric with a skewness coefficient of −2.55. Market-makers have a 5% risk of losing 21 million dollars or more on any given day.

The fact that market-makers earn a profit on average from delta-hedging their inventory generates the positive trend observed in the middle panel of Figure 3. In aggregate, market-makers earn 9 million dollars monthly from their delta-hedged positions.

Finally, comparing the top panel with the bottom panel reveals the substantial impact market variance risk has on dealers' wealth. Unlike equity stock market-makers, a large portion of changes in option market-makers' wealth is driven by the fluctuations in their delta-hedged inventory. In our sample, the absolute value of $P\&L_t$ is on average 1.5 times bigger than BA_t.

So far, we have established that the representative SPX market-maker faces substantial variance risks from carrying large inventories over time. In option markets, dealers often trade among themselves in order to manage part of the risks they face. However, whenever end-users have large net exposure to market variance, so do SPX market-makers in the aggregate. Because of market-makers large exposure to market variance, part of their required compensation should be embedded in the variance risk premium.

[7] This fee includes \$0.33 charged by the CBOE and \$0.03 charged by the Option Clearing Corporation for clearing cost.

2.3 Methodology and Predictions

Our empirical analysis is based on two different specifications for the variance risk premium. Our first serie of tests is conducted on Carr and Wu (2009) log-variance risk premium, $LogVRP_{t,T} \equiv \ln(RV_{t,T}/RNV_{t,T})$.[8] We then investigate the robustness of our results using $VRP_{t,T}$ as defined by (2.1).

The methodology we follow is adapted from Bollen and Whaley (2004). More precisely, we regress the daily changes in the variance risk premium against the explanatory variables and lagged changes. When working with the log-specification, we run

$$\begin{aligned}\Delta LogVRP_{t,T} = Intercept + \beta_1^{Inv} InvRisk_{t-1} + \beta_2^{Inv} \left(\Delta W_t \cdot InvRisk_{t-1}\right) \\ +\beta^c Control_t + \beta^{VRP} \Delta LogVRP_{t-1,T} + \varepsilon_t,\end{aligned} \tag{2.8}$$

where $\Delta LogVRP_{t,T} \equiv LogVRP_{t,T} - LogVRP_{t-1,T}$, and similarly for $\Delta VRP_{t,T}$.

If the variance risk premium captures part of dealers' required compensation, we should expect a positive relationship between lagged inventory risk and changes in the variance risk premium. The more negative (positive) inventory risk is the more negative (positive) the variance risk premium should be as such a relationship implies positive returns for dealers. Therefore, lagged inventory risk should positively impact changes in variance risk premium. Our prediction is thus $\beta_1^{Inv} > 0$.

As their wealth decreases, dealers should command a higher compensation. After controlling for their lagged exposure, interacting cotemporaneous changes in wealth with lagged inventory risk should capture the way dealers dynamically update their required compensation as their wealth fluctuates. When market-makers are experiencing losses and are negatively exposed to market variance, $\Delta W_t \cdot InvRisk_{t-1}$ is positive. Given the sign of market-makers' exposure, a higher compensation implies a decrease in the variance risk premium. Therefore, the interaction of changes in wealth with lagged inventory should be negatively related to changes in the variance risk premium. Our prediction is $\beta_2^{Inv} < 0$.[9]

Motivated by previous studies, we include a series of cotemporaneous control variables. We discuss them in the next section.

2.4 Control Variables

Carr and Wu (2009) show that part of the variation in the variance risk premium is related to market index returns. To control for the variation in variance risk premium induced by contemporaneous

[8] Because the distributions of the two measures of variance are positively skewed, the log-specification alleviates the effect of extreme values.

[9] Note that a similar argument also predicts $\beta_2^{Inv} < 0$ when inventory risk is positive.

changes in market conditions, we calculate the log-returns of the S&P 500 index and denote this variable $S\&P500LogRet_t$.

Egloff, Leippold, and Wu (2010), Todorov (2010), and Aït-Sahalia, Karaman, and Mancini (2014) among others study the impact of jumps on the variance risk premium. We follow Cremers, Halling, and Weinbaum (2014) and construct an aggregate jump factor. On each day, we calculate the returns on two zero-beta at-the-money SPX straddles of maturities T_1 and T_2 with $T_1 < T_2$. We denote by r_t^{S1} and r_t^{S2} the returns on the two straddles.[10] These daily returns are then combined such that

$$JumpFactor_t \equiv r_t^{S1} - \left(\frac{Vega_t^{S1}}{Vega_t^{S2}}\right) \cdot r_t^{S2}, \tag{2.9}$$

where $Vega_t^{S1}$ and $Vega_t^{S2}$ denote the vega of each straddle. By construction $JumpFactor_t$ has zero delta, zero vega, and positive gamma, and thus captures the large fluctuations in the S&P500 index.[11]

Bollen and Whaley (2004) document the effect of net buying pressures on option implied volatility. Through their impact on implied-volatility, net buying pressures potentially influence the variance risk premium. To disentangle the effect of inventory risk relative to net buying pressures on the variance risk premium, we calculate on each day Bollen and Whaley's net buying pressure variable

$$NetByingPressure_t \equiv \sum_j \left(BO_t^j - SO_t^j\right) \cdot \frac{abs\left(\Delta_t^j\right)}{Volume_t}, \tag{2.10}$$

where $Volume_t$ is the aggregate volume of SPX options, and $abs(.)$ denotes the absolute value.

Buraschi, Trojani, and Vedolin (2014) establish that investors' disagreement influences the variance risk premium. Empirically, dispersion of analyst forecasts is often used to gauge investors' disagreement. However, data limitation restricts the daily calculation of this measure. Because a higher disagreement increases trading motives among investors, we rely on S&P500 index volume data and compute unexpected changes in index trading volume. For every day we calculate the difference between S&P 500 index volume on that day, and the average volume of the last 90 trading days. We denote this variable $Disagreement_t$.

2.5 Time Variation in the Variance Risk Premium

Table 3 presents the estimated coefficients for regressing the daily changes in log-variance risk premium on inventory risk and market-makers' wealth. We also include the control variables previously

[10] When constructing the first straddle, we choose T_1 to be in the month following the current month, and at least fifteen days to maturity. Similarly, T_2 is chosen to be in the month following T_1.

[11] The gamma of an option is the second derivative of the option price with respect to the underlying asset price.

discussed. We report the results for the full sample in regressions (1) and (2). Moreover, we divide the sample in two sub-samples of similar length 1997-2004 and 2005-2011. This allows us to assess the importance of the financial crisis in influencing the results obtained. Note that each variable is standardized to have unit variance in order to facilitate the interpretation of the coefficients. For each parameter, the p-value is computed using Newey-West p-value with 8 lags to capture autocorrelation in the residuals.

Regressions (2), (4), and (6) in Table 3 establish the importance of inventory risk and market-maker wealth for explaining the variance risk premium. Both variables are statistically significant with their expected sign. For the two subsamples, their inclusion increases the adjusted R-Squared by 9% relative to regressions (3) and (5). Given an average adjusted R-Squared of 42% this corresponds to a $0.09/0.42 = 21\%$ increase in explanatory power.

Interestingly, it is for the sample period including the financial crisis that the impact of inventory risk is the greatest. A one standard deviation decrease in inventory risk causes a 1.2% decrease in the variance risk premium. When inventory risk equals its 2005-2011 sample average, a one standard deviation decrease in market-maker wealth is associated with a 4.55% decrease in the variance risk premium. The effect of inventory is magnified when market-makers experience dramatic losses. Conditioning on a market-makers' loss at its 90^{th} percentile (i.e. $\Delta W_t =$-$72,898,000$ million dollars), a one standard deviation decrease in inventory risk results in a 9% decrease in next day's log-variance risk premium.

The relation between S&P 500 log-returns and the variance risk premium is strongly significant. The variance risk premium tends to decrease when index returns decrease. In regressions (1), (3), and (5) in Table 3, S&P 500 log-returns and lagged changes in variance risk premium jointly account for 80% of the R-Squared documented.[12]

Note the way aggregate jumps are negatively related to changes in variance risk premium. By construction, the returns of the jump factor are high whenever the S&P 500 abruptly decreases. Thus, the estimated coefficient suggests that large negative jumps result in a more negative variance risk premium. This is consistent with Todorov (2010) who studies the impact of market jumps on the variance risk premium.

Bollen and Whaley (2004) demonstrate that net buying pressures increase implied volatility. Through their effect on implied volatility, high buying pressures should result in a lower variance risk premium. From Table 3, we see that net buying pressure is negatively related to the variance risk premium. However, the variable is only significant in regression (5).

Consistent with Buraschi, Trojani, and Vedolin (2014), the loadings estimated for disagreement

[12] Bivariate regressions of changes in log-variance risk premium on index log-returns and lagged changes have an adjusted R-Squared of 30% on average for the full sample and subsample tests.

are consistently negative. When investors' disagreement increases, the difference between realized and implied volatility tends to become more negative.

The Term Structure of Variance Risk Premia

The term structure of the variance risk premium is a topic of recent interest. It has been studied in Egloff, Leippold, and Wu (2010) and Aït-Sahalia, Karaman, and Mancini (2014) among others. To quantify the impact each variable has on the term structure of variance risk premia, we construct measures of the variance risk premium for various horizons. In addition to the one-month horizon, we focus on four additional horizons of 60, 90, 180, and 270 days respectively. We present the full sample results in Table 4.

Table 4 reports the result obtained for the log-variance risk premium. Several interesting findings emerge. Note that the effect of inventory risk and fluctuations in market-makers wealth are robust across horizons. Interestingly, there is a term structure of the coefficients estimated. For most variables, the magnitude of their influence on the variance risk premium decreases as the horizon increases. Accordingly, the effect of inventory risk and market-makers wealth is most prominent for short-term variance risk premia. This result complements the finding of Aït-Sahalia, Karaman, and Mancini (2014), who conclude that investors' fear of a market crash is mostly captured in short-term variance risk premia. For horizons of at least sixty days, inventory risk and market-makers' wealth have the second greatest economic impact after the index returns, and before aggregate jumps.

Forecasting Market Variance

So far, we have used future realized variances to proxy expected physical variances. We now investigate the robustness of our results when expected physical variance is set to the one-step forecast of predictive model.

We adopt a HAR-RV dynamic in the spirit of Corsi (2009) in order to model realized variances over time. Using rolling-windows of 252 observations we estimate the model for each maturity on each day. We then set $RV_{t,T}$ to the one-step-ahead model forecast. The Appendix A provides details on the methodology adopted. We present the average of the daily parameter estimates, p-values, and R-Squared from fitting the HAR-RV model in Panel A of Table 5. The high p-values demonstrate the difficulty to precisely estimate each of the coefficient. Nevertheless, the high R-Square demonstrates the model ability to forecast future market variance movements. Using the model prediction for $RV_{t,T}$ we construct measures of the variance risk premium for each horizon. We refer to Panel B for descriptive statistics on the variance risk premium implied by the HAR-RV dynamic.

Based on the variance risk premia constructed we regress the changes of log-variance risk premium on the explanatory variables for each horizon. Details on the regressions are provided in Table A.1 in the Online Appendix. The adjusted R-Squared ranges from 40% to 49%. More importantly, the effect of inventory risk and market-makers' wealth are robust to the use of the HAR-RV model.

In Table 6, we present the average coefficients, p-values, and adjusted R-Squared across horizons. In the first column, the values reported are based on the results from Table 4 while the second column reports the average measures infer from Table A.1. As readily apparent, the coefficients are close irrespectively of the way expected physical variances are measured.

Additional Robustness Tests

In Table 7, we assess the robustness of our empirical analysis when the variance risk premium is measured as $RV_{t,T} - RNV_{t,T}$ where $RV_{t,T}$ is proxy by future realized variances as in Table 4. The impact of inventory risk and fluctuations in market-makers wealth remain robust. Within the sample, univariate regressions of $\Delta LogVRP_{t,T}$ on $\Delta VRP_{t,T}$ have on average a factor loading of 10 across horizons. Therefore, multiplying the parameter estimates in Table 7 by 10 renders the comparison with Table 4 possible. With this adjustment, the coefficient at the one-month horizon becomes 0.60 for inventory risk and -1.7 for the interaction variable. Once adjusted, the parameter values in Table 7 relate well to the one reported in Table 4.

Recall that our calculation of inventory delta-hedged profits and losses relies on option mid-quotes. Because market-makers carry net short positions on average, intermediaries trade at the ask price more often than they trade at the bid price. We now assess if this impacts the results documented. To do so, we use end-of-day ask prices to calculate intermediaries' delta-hedged profits and losses on each day. Based on these daily estimates we compute the daily changes in market-makers wealth and proceed with the regression analysis. Details of the regression is available in Table A.2. Across horizons, the coefficient for the interaction of inventory risk with changes in wealth is -2.15 on average. This is close to the average estimate of -2.13 obtained when using mid-quotes. We conclude that the coefficient estimates are robust to the measurement of market-makers' wealth.

2.6 The Financial Crisis

Typically, the index option market functions similarly to an insurance market. A clientele of institutional investors predominantly buys SPX options, which causes market-makers' inventory risk to be negative on average. On November 20, 2008, the VIX Index reached a high of 80.86%. Around the same time, end-users were heavily shorting SPX options. These large net sell orders resulted in an increase in inventory risk of more than 1.8 billion dollars during the month of November. By November 20, CBOE market-makers carried more than 2.5 billion dollars positive exposure to market volatility. This risk exposure represents more than 8% of the total capitalization of SPX options at that time.

To understand the risk and reward incurred from being positively exposed to market volatility during the financial crisis, Table 8 presents the return statistics of delta-hedged near-the-money options. Because these options are close to the money, they are highly sensitive to market variance. For comparison purposes, we report the statistics for the full sample and the financial crisis separately.

While losing most of the time, the delta-hedged positions earn 6.26% monthly during the crisis period. Thus, around the same time market-makers' exposure to market variance is positive, long positions in near-the-money options are profitable on average. Note however the tremendous risk exposure from these options. During the financial crisis, the volatility of daily returns peaks at 95%, and its excess kurtosis is about 11. This result further stresses the critical role of index option market-makers in undertaking large risk to allow end-users to hedge against and speculate on market volatility.

Overall, the evidence presented so far supports the idea that part of the variance risk premium captures index option market-makers' compensation for their exposure to market variance. Motivated by these findings, we now develop a theoretical model in which the physical market variance is dynamic and a risk-averse representative market-maker endogenously quotes index options, and thus influences the variance risk premium.

3 The Model

The economy is cast in a continuous time framework in which the underlying source of uncertainty is driven by two independent Brownian motions Z^S and Z^V.[13] Market participants have a finite investment horizon T, and can invest in the market index S_t, which evolves according to

$$\frac{dS_t}{S_t} = (\mu - q)\, dt + \sqrt{V_t}\left(\sqrt{1-\rho_V^2} dZ_t^S + \rho_V dZ_t^V\right), \quad S_0 \;\; \text{known}, \tag{3.1}$$

where μ is the market premium, q is the dividend yield, and V_t is the market variance which follows a CEV dynamic

$$dV_t = \kappa(\theta - V_t)dt + \delta V_t^{\eta} dZ_t^V, \quad V_0 \;\; \text{known}, \tag{3.2}$$

[13] The information available to agents consists of the trajectories generated by the two Brownian motions (Brownian filtration $\mathcal{F}$). The underlying probability space is $(\Omega, \mathcal{F}, P)$, P being the physical probability measure.

where θ denotes the unconditional variance, κ is the speed of mean reversion, δ is the volatility of volatility, and η determines the variance elasticity.[14] In (3.1), ρ_V captures the correlation between market return and market variance innovations. In addition to the market index, market participants can invest in a risk-free bond

$$\frac{dB_t}{B_t} = rdt, \quad B_0 = 1, \tag{3.3}$$

with constant interest rate r. The economy is endowed with a stochastic discount factor (SDF) which reflects aggregate preferences. As in the portfolio literature (see Detemple, Garcia, and Rinsdisbacher, 2003, 2005; Detemple and Rinsdisbacher, 2010; and Elkamhi and Stefanova, 2011), the form of the SDF is exogenously given in the model. This SDF follows

$$\frac{d\xi_t}{\xi_t} = -rdt - \phi_t^S dZ_t^S - \phi_t^V dZ_t^V, \quad \xi_0 = 1, \tag{3.4}$$

where ϕ_t^S and ϕ_t^V are the market prices of risks. The (instantaneous) variance risk premium is the product of the market price of variance risk with the quantity of variance risk, that is $VRP_t = \phi_t^V \cdot (\delta V_t^\eta)$. Therefore, the variance risk premium and the price of variance risk are substitutes as ϕ_t^V can be replaced by $VRP_t/(\delta V_t^\eta)$ in (3.4).

The variance risk premium is to be determined through the trading activity in index options. We denote by f_t^j European index calls and puts where j identifies a particular option. In the option market, two types of agents interact. End-users have an exogenous need to get exposure to index options. We denote by $\Phi_t^{EU,j}$ end-users' net demand for option j. To meet this demand, a representative market-maker provides liquidity for index options. Since the physical dynamics of the market index, the market variance and the SDF are all exogenous, the demand and supply for index options only influences the determination of ϕ_t^V and VRP_t within our framework.

In the next proposition, we present the pricing rule adopted by the market-maker to quote each option.

Proposition 1 *Given (3.1), (3.2), and (3.4) applying Ito's lemma to f_t^j implies the following dynamic for the price of option j under the P-measure*

$$\begin{aligned} df_t^j &= d\Delta Rep_t^j + \vartheta_t^j \cdot dF_t^V \\ &= d\Delta Rep_t^j + \vartheta_t^j \cdot \left(VRP_t dt + \delta V_t^\eta dZ_t^V\right), \end{aligned} \tag{3.5}$$

[14] Equation (3.2) nests many stochastic volatility models. For instance, Heston (1993) is obtained when $\eta = 1/2$, and Jones (2003) discusses the model for $\eta > 1$.

where ΔRep_t^j *corresponds to the delta replication of* f_t^j, $\vartheta_t^j \equiv Vega_t^j / \left(2\sqrt{V_t}\right)$ *is the sensitivity of option* j *to the market variance risk factor* F_t^V, $VRP_t \equiv \frac{1}{dt}\left(E_t^P[V_{t+dt}] - E_t^Q[V_{t+dt}]\right) = (\delta V_t^\eta) \cdot \phi_t^V$ *is the (instantaneous) variance risk premium, and* $\delta V_t^\eta dZ_t^V$ *is the aggregate variance risk.*

Proof. See Appendix B. ∎

Under the Black and Scholes (1973) assumptions, index options can be perfectly replicated by holding the appropriate amount of the market index and the risk-free bond. When the market variance is stochastic, perfect replication is no longer achievable through the trading of S_t and B_t only. As a result, the price dynamic of index options can be decomposed into two components. As in Black and Scholes (1973), the first component, denoted $d\Delta Rep_t^j$, corresponds to the delta replication of f_t^j. In addition to $d\Delta Rep_t^j$, the entire cross-section of index options is affected by the market variance risk factor.

When $\Phi_t^{MM,j}$ denotes the market-maker's inventory of option j, the market clearing condition for index options is

$$\Phi_t^{MM,j} + \Phi_t^{EU,j} = 0 \text{ for all } j \Rightarrow InvRisk_t = \sum_j \Phi_t^{M,j} Vega_t^j = -\sum_j \Phi_t^{EU,j} Vega_t^j. \tag{3.6}$$

Whenever end-users' exposure to market volatility, $\sum_j \Phi_t^{EU,j} Vega_t^j$, does not cancel out across index options, then neither does market-maker's inventory risk. Consequently, the representative market-maker will be non-trivially dynamically exposed to the market variance risk factor.

For tractability reasons, we do not endogenize end-users' trading motives for index options. Instead, we build on the work of Amihud and Mendelson (1980) and Gârleanu, Pedersen, and Poteshman (2009) among others and model the fluctuations of inventory risk exogenously. As apparent in Figure 2, the time-series of inventory risk shares feature with volatility (i.e. clustering, autocorrelation, and reversal). This is consistent with the market microstructure literature who find that stock market-makers' inventory mean revert (see Madhavan and Sofianos, 1998, Hansch, Naik, and Viswanathan, 1998, and Naik and Yadav, 2003). To capture these statistical properties, we define the following mean-reverting dynamic for inventory risk

$$dInvRisk_t = \lambda(\alpha - InvRisk_t)dt + \psi V_t dt + \sigma InvRisk_t \left(\sqrt{1-\rho_{Inv}^2}dZ_t^S + \rho_{Inv} dZ_t^V\right), \tag{3.7}$$

where $InvRisk_t$ satisfies (3.6). In the dynamic above λ captures the speed of mean reversion, α influences the level of inventory risk, and σ is the volatility parameter. Arguably, fluctuations in market variance should influence the aggregate net demand for index options and thus inventory risk. To account for this, we allow the dynamic (3.7) to depend on V_t and dZ_t^V. In the previous equation, ψ captures the sensitivity of inventory risk to market variance, and ρ_{Inv} measures its correlation with market variance innovations.

We determine the variance risk premium through the maximization of the market-maker's expected utility of terminal wealth. For a given admissible investment strategy $\tilde{\pi}_t \equiv \left(\pi_t^B, \pi_t^S, \{\pi_t^{f,j}\}\right)$, the market-maker's self-financing wealth dynamic is

$$\frac{dW_t}{W_t} = \pi_t^B \cdot \frac{dB_t}{B_t} + \pi_t^S \cdot \left(\frac{dS_t}{S_t} + q dt\right) + \sum_j \pi_t^{f,j} \cdot \frac{df_t^j}{f_t^j}, \quad \text{with} \quad W_0 = w, \tag{3.8}$$

where each π is expressed in percentage of wealth, the term qdt accounts for the reinvestment of dividends, and w denotes his initial endowment. The problem faced by the market-maker can be written as

$$\max_{\tilde{\pi}_t} \; E^P[U(W_T)] \quad \text{subject to} \quad \begin{array}{l} \cdot \; (3.7) \\ \cdot \; (3.8) \text{ with } W_t \geqslant 0\text{: } t \in [0, T], \end{array} \tag{3.9}$$

where $U(.)$ is the market-maker's utility function. In the model, the representative market-maker determines his trading strategy given prices. Our goal is to invert this mapping and infer the variance risk premium in (3.5) that induces the market-maker to clear the index option market.

4 Model Implications

4.1 The Structure of the Variance Risk Premium

To understand the way the variance risk premium depends on inventory risk and the market-maker's wealth, the following proposition is needed.

Proposition 2 *At time $t \in [0, T]$, if the market-maker is myopic with $U(W_t) = \ln(W_t)$, the variance risk premium is given by*

$$VRP_t = \rho_V \delta V_t^{\eta} \left(Sharpe_t\right) + 0.5(1 - \rho_V^2)\delta^2 V_t^{2\eta - 0.5} \left(\frac{InvRisk_t}{W_t}\right), \tag{4.1}$$

where $Sharpe_t \equiv \frac{\mu - r}{\sqrt{V_t}}$ is the market Sharpe ratio, $InvRisk_t = -\sum_j \Phi_t^{EU,j} Vega_t^j$ captures the sensitivity of the market-maker's inventory to the variance risk factor F_t^V defined in Proposition 1, and W_t is the market-maker wealth.

Proof. *See Appendix C.* ■

This proposition provides several insights. The decomposition in (4.1) separates the variance risk premium into two components. The first term is a function of the market Sharpe ratio. Innovations in market variance are correlated with index returns. Consequently, the variance risk premium inherits of the market Sharpe ratio. The greater the $abs(\rho_V)$ the higher the dependence of the variance risk premium on $Sharpe_t$. Arguably, index returns are influenced by the realisations of the market Sharpe ratio. Therefore, equation (4.1) helps understand the strong statistical significance found when regressing the changes in variance risk premium measures on $S\&P$ 500 log-returns.

Whenever $abs(\rho_V) < 1$, index options cannot be perfectly hedged through the trading of S_t and B_t. Consequently, the market is incomplete from the market-maker's perspective. In this case, the market-maker requires an additional compensation relative to $\rho_V \delta V_t^{\eta} \left(Sharpe_t\right)$. This additional premium is proportional to the ratio of inventory risk to wealth. Since $(1-\rho_V^2)\delta^2 V_t^{2\eta-0.5} \geqslant 0$ and $W_t \geqslant 0$, the variance risk premium is positively influenced by inventory risk. The more negative is the exposure of market-maker to market variance, the lower the variance risk premium will be. Consequently, the model captures the positive sign of the coefficient obtained when regressing the changes in variance risk premium against market-makers' exposure to market variance.

Several studies have shown that the premium for market variance risk is negative on average (see, among others, Bakshi and Kapadia, 2003, Vilkov, 2008, and Carr and Wu, 2009). Given (4.1), the variance risk premium is negative whenever

$$\frac{InvRisk_t}{W_t} < \frac{-\rho_V \left(Sharpe_t\right)}{0.5(1-\rho_V^2)\delta V_t^{\eta-0.5}}. \tag{4.2}$$

For the empirically relevant case $\rho_V < 0$ and $Sharpe_t > 0$, a sufficient condition for the variance risk premium to be negative is a negative inventory risk. Since index option market-makers typically have a negative exposure to market variance, the model offers an interesting channel for explaining the negative variance risk premium found in previous studies.

As apparent in the middle and bottom panels of Figure 1, the variance risk premium is occasionally positive. When inequality (4.2) is not satisfied, and $\rho_V < 0$ and $Sharpe_t > 0$, a positive inventory risk exposure results in a positive variance risk premium. Hence, the model can also explain a positive variance risk premium at times when option market-makers are positively exposed to market variance.

Finally, note that a positive inventory risk will not result in a positive variance risk premium as long as (4.2) is satisfied.

4.2 Optimal Market-Maker Wealth

Most classical inventory models assume that dealers have access to unlimited capital (see, among others, Holl and Stoll, 1981 and 1983, and Mildenstein, and Schleef, 1983). Recently, Gromb and

Vayanos (2002) and Brunnermeier and Pedersen (2009) relax this assumption. In Brunnermeier and Pedersen (2009) market-makers' limited funding capacity influence their liquidity provisions. When market-makers' margin requirements approach their available capital, intermediaries provide less liquidity which leads to price distortions. Similar predictions are obtained by Gromb and Vayanos (2002). In our model, the intermediary's financial constraint has also important pricing implications.

In Proposition 2, inventory risk is normalized by the market-maker's wealth. Consequently, inventory risk will matter the most for the variance risk premium when the market-maker is poor (i.e. for low values of W_t). In contrast, the effect of inventory risk vanishes when the market-maker is unconstrained, that is, when his wealth goes to infinity. This result justifies the strong significance found in all the empirical tests for the interaction of inventory risk with changes in market-makers' wealth.

In the model, the market-maker's wealth is endogenous and if $U(W_t) = \ln(W_t)$ then we have

$$W_t = 1/\left(\gamma \xi_t\right), \tag{4.3}$$

where γ is the shadow price of the market-maker's budget constraint. Since the market-maker's marginal utility is strictly increasing, γ is uniquely defined by $W_0 = E^P[\xi_T W_T] = 1/\gamma \Rightarrow W_t = W_0/\xi_t$. Consequently, at any point of time, the market-maker's wealth is proportionally related to the ratio of his initial endowment W_0 to the SDF. Given (3.4), we can apply Ito's lemma to W_t to obtain the dynamic of the market-maker's optimal wealth

$$dW_t = \left(r + \left(\phi_t^S\right)^2 + \left(\phi_t^V\right)^2\right) W_t dt + \phi_t^S W_t dZ_t^S + \phi_t^V W_t dZ_t^V. \tag{4.4}$$

We see that the market-maker's wealth is driven by the two aggregate shocks and their prices of risk. As we will discuss it further in the next section, the dependence of the market-maker's wealth on the price of variance risk has interesting implications for the risk-neutral distribution of the market return.

4.3 Risk-Neutral Dynamics

To discount the option payoff at maturity using the risk-free rate, all underlying processes must be risk-neutralized. In light of this, the result in the next proposition is particularly relevant for pricing.

Proposition 3 *When the SDF follows (3.4), the processes describing the evolution of the risk-*

neutral economy are

$$dS_t = (r-q)\,S_t dt + \sqrt{V_t} S_t \left(\sqrt{1-\rho_V^2} d\tilde{Z}_t^S + \rho_V d\tilde{Z}_t^V\right) \tag{4.5}$$

$$dV_t = \kappa(\theta - V_t)dt - VRP_t dt + \delta V_t^{\eta} d\tilde{Z}_t^V \tag{4.6}$$

$$\begin{aligned} dInvRisk_t &= \lambda(\alpha - InvRisk_t)dt + \psi V_t dt + \sigma InvRisk_t \left(\sqrt{1-\rho_{Inv}^2} d\tilde{Z}_t^S + \rho_{Inv} d\tilde{Z}_t^V\right) \\ &\quad -\sigma InvRisk_t \left(\sqrt{1-\rho_{Inv}^2}\phi_t^S + \rho_{Inv}\phi_t^V\right) dt \end{aligned} \tag{4.7}$$

$$dW_t = rW_t dt + \phi_t^S W_t d\tilde{Z}_t^S + \phi_t^V W_t d\tilde{Z}_t^V. \tag{4.8}$$

where $\tilde{Z}_t^S$ and $\tilde{Z}_t^V$ are risk-neutral, $\phi_t^S = Sharpe_t/\sqrt{1-\rho_V^2} - \rho_V \phi_t^V/\sqrt{1-\rho_V^2}$ is obtained by imposing the index no-arbitrage condition, $\phi_t^V = VRP_t/(\delta V_t^{\eta})$, and VRP_t satisfies Proposition 2.

Proof. *The Girsanov theorem implies that $dZ_t^S = d\tilde{Z}_t^S - \phi_t^S dt$ and $dZ_t^V = d\tilde{Z}_t^V - \phi_t^V dt$. Using this result in (3.1), (3.2), (3.7), and (4.4) delivers the risk-neutral processes.* ■

When end-users' demand for index options increases, it decreases inventory risk. A decrease in inventory risk implies a more negative variance risk premium on average. From (4.6), we see that changes in the risk-neutral market variance are negatively influenced by the variance risk premium. Thus, an increase in end-users' demand will result in a higher risk-neutral variance. This model prediction can be related to Bollen and Whaley (2004) who document that changes in the implied volatility of OTM index puts are positively influenced by end-users' net buying pressure. However, our model suggests that only the variance exposure of market-maker's total risk exposure should impact changes in risk-neutral volatility.

The nonlinearities in the joint dynamic of risk-neutral variance, inventory risk and market-maker's wealth have interesting implications for market risk-neutral skewness. When variance risk is high (i.e. $d\tilde{Z}_t^V > 0$) market returns are low when $\rho_V < 0$. At times when inventory risk and the variance risk premium are both negative high variance risk reduces the market-maker's wealth since $\phi_t^V d\tilde{Z}_t^V < 0$ in (4.8). A lower wealth implies a more negative ratio of inventory risk to wealth which further decreases the variance risk premium. This feedback effect between the variance risk premium and the ratio of inventory risk to wealth amplifies the high level of risk-neutral variance in bad times. This allows the model to generate substantial negative skewness which is appealing given the challenge of standard SV models to explain the cross-section of out-of-the-money puts.

Overall, the predictions delivered by the model are in line with empirical stylized facts. In the next section, we estimate the model and assess the role of inventory risk and market-maker wealth for the valuation of index options.

5 Estimation and Fit

In this section, we first describe our estimation methodology. Subsequently, we report on parameter estimates and model fit. Finally we relate the estimated parameters to the economic impact of inventory risk and market-maker's wealth on index option prices.

5.1 Estimation Methodology

Several approaches have been proposed in the literature for estimating stochastic volatility models. Aït-Sahalia and Kimmel (2007) and Jones (2003) estimate volatility model parameters using bivariate time-series of returns and at-the-money implied volatility. Pan (2002) uses GMM to estimate the objective and risk-neutral from returns and option prices. Christoffersen, Jacobs, and Mimouni (2010) adopt a particle filtering approach to estimate various alternatives to Heston (1993) based on returns and a large panel of option prices combined.

For estimating the model, we adopt a two steps procedure. In a first step, we filter the physical parameters of the market variance dynamic (3.2) along with the spot variances V_t on index returns. Taking the filtered spot variances and the objective variance parameters as given, we estimate the dynamic of inventory risk and market-maker's wealth using a large panel of SPX put prices. Both $InvRisk_t$, and W_t are latent variables in the model. However, to avoid overfitting the data we constrain $InvRisk_t$ to be equal to its empirical value (2.5). In addition, we fix the market-maker's initial wealth to $W_0 = w$ at the beginning of every day. Based on these initial values and on the filtered V_t, we exploit the results in Proposition 2 and 3 to simulate the economy forward and infer the dynamics (4.7) and (4.8) consistent with observed index option prices.

For estimation purposes, we set the time-step Δt to 1/365 and the expected return net of dividend yield $\mu - q$ to its sample average $\frac{1}{T-1}\sum_{t=2}^{T}\left(S_t - S_{t-1}\right)/S_{t-1} \times 365$, where T denotes the last day in the sample.

We now describe each step in more details.

Step 1: Filtering the Variance Dynamic Using S&P 500 Returns

We need to estimate the structural parameters $\Theta^V \equiv \{\kappa, \theta, \delta, \rho_V, \eta\}$ in (3.2) along with the vector of spot variances $\{V_t\}_{t=1,2,\ldots,T}$. To this end, we adopt the particle filtering algorithm (PF henceforth). The PF offers a convenient way to estimate stochastic volatility models. It has recently been applied in Johannes, Polson, and Stroud (2009), Christoffersen, Jacobs, and Mimouni (2010), and Malik and Pitt (2011) among others.

Let $\left\{V_t^j\right\}_{j=1}^N$ denotes the smooth resampled particles where N defines the number of particles which is set to $10,000$. Using the algorithm described in the Appendix D, we estimate on each day the likelihood of observing S_{t+1} given V_t^j and S_t, and denote it $\tilde{P}_t^j\left(V_t^j,\Theta^V\right)$. Based on the likelihood of each particle calculated on each day, the MLIS criterion applied to estimate Θ^V is

$$\hat{\Theta}^V = \arg\max \sum_{t=1}^{T} \mathcal{L}_t, \tag{5.1}$$

where $\mathcal{L}_t \equiv \ln\left(\frac{1}{N}\sum_{j=1}^{N}\tilde{P}_t^j\left(V_t^j,\Theta^V\right)\right)$ is the model daily log-likelihood. On day t, the filtered spot variance is obtained by averaging the smooth particles

$$\hat{V}_t = \frac{1}{N}\sum_{j=1}^{N} V_t^j. \tag{5.2}$$

Next, we discuss the estimation of inventory risk and market-maker's wealth based on SPX options.

Step 2: Estimating Inventory Risk and the Market-Maker's Wealth

The model does not allow for closed-form solution for option prices. Consequently, we rely on Monte-Carlo methods for estimating the inventory risk and market-maker's wealth dynamic embedded in SPX options. Taking $\hat{\Theta}^V$ and $\left\{\hat{V}_t, Inv\hat{R}isk_t\right\}_{t=1,2,\ldots,T}$ as given, where $Inv\hat{R}isk_t$ corresponds to (2.5), we estimate $\Theta^{Inv} \equiv \{\lambda,\alpha,\psi,\sigma,\rho_{Inv}\}$ and w by minimizing the sum of implied volatility squared errors (SIVSE)

$$\left\{\hat{\Theta}^{Inv},\hat{w}\right\} = \arg\min \sum_{j,t}^{N_J}\left(IV_{j,t} - IV_{j,t}^M\left(\Theta^{Inv}, w, \hat{\Theta}^V, \hat{V}_t, Inv\hat{R}isk_t\right)\right)^2, \tag{5.3}$$

where N_J is the total number of observations, $IV_{j,t}$ is the option j implied volatility on day t, and $IV_{j,t}^M$ is the model implied volatility. We use Black-Scholes to calculate implied volatilities for both market and model prices. When calculating model prices, we follow the algorithm described in the Appendix E using $10,000$ Monte-Carlo paths.

5.2 Parameter Estimates

Market Variance Parameters

In Panel A of Table 9, we present descriptive statistics on the daily S&P 500 returns. In-sample, the average market return net of dividend yield is 5.86%. This low average return estimate is partly influenced by the abrupt drop in the S&P 500 index during the crisis period (see Figure 1). During the 1997-2011 period, the sample variance is 4.58% annually, which corresponds to a 21% average volatility.

In Panel B, we report the estimated parameters for the CEV dynamic (3.2). The sample MLIS is $11,746$. The estimated θ is close to the sample variance in Panel A. Moreover, ρ_V is large and negative. In the model, a large and negative ρ_V is important to generate enough negative skewness in the S&P 500 return distribution. Large fluctuations of market variance (i.e. high volatility of dV_t) helps the model generates additional variance in the return process. This is required to capture the kurtosis of the S&P 500 return distribution. Given the parameter estimates, the variance of dV_t is on average equal to $d\langle V,V\rangle_{V=\theta} = \delta\theta^{\eta}dt = 0.06dt$. Thus, the high volatility of volatility δ is required by the model to generate enough fluctuations of variance given $\eta = 0.90$. Finally, note how the index data require a slow variance mean-reversion speed. The coefficient of 2.91 corresponds to a daily variance persistence of $1 - \kappa/365 = 0.99$.

For comparison purposes, we also report in Panel C the parameters obtained for the Heston model with $\eta = 1/2$. The model MLIS is close to the likelihood obtained by the CEV dynamic. However, the two models require different structural parameters to explain the data. For instance, note the differences of mean reversion speed and unconditional variance between the two models. The Heston dynamic requires a higher speed reversion but has a lower long-term variance. Moreover, the two models also require different volatility of volatility. This is partly driven by the difference in η between the two models. Because $\sqrt{\theta} = \sqrt{0.0408} = 0.20$ for Heston (1993) is substantially higher than $\theta^{\eta} = 0.0458^{0.90} = 0.06$ for CEV, the Heston model requires a smaller volatility of volatility parameter than CEV to fit the S&P 500 return distribution. For Heston (1993), we have on average $d\langle V,V\rangle_{V=\theta} = \delta\sqrt{\theta}dt = 0.04dt$ which is slightly less than the CEV model.

In Figure 4, we plot the time-series of filtered spot volatilities $\sqrt{\hat{V}_t}$ annualized and expressed in percentage for both models. As expected, the time-series of physical spot volatilities share common features with the VIX Index in Figure 2. Comparing both panels we see that the filtered spot volatilities display similar patterns most of the time. However, the filtered spot variance for the CEV model are substantially higher than the one filtered for Heston (1993) during the crisis. Arguably, this is due to the higher elasticity of the CEV model which requires a higher level of spot volatilities to generate sufficient kurtosis during the crisis period.

Inventory Risk and Market-Maker's Wealth Coefficients

When calibrating inventory risk and market-maker wealth, we rely on OptionMetrics' volatility surface data. In order to speed up estimation, we restrict attention to puts observed on the first Wednesday of each month with moneyness between 0.9 and 1.1, and with 2, 3, and 6 months to maturity. The final option sample consists of $6,292$ put observations during the 1997-2011 sample period.

In Panel A of Table 10, we report the estimated coefficients for $\hat{\Theta}^{Inv}$ obtained from minimizing the sum of squared errors (2.5). The mean reversion parameter is 10.73. A high mean reversion speed is necessary to explain the abrupt reversal of inventory risk observed during the crisis period in the bottom panel of Figure 2. The daily inventory persistence is $1 - \lambda/365 = 0.97$. Surprisingly, the persistence implied by options is close to the persistence of 0.98 obtained when fitting an AR(1) model on the spot inventory risk measures. The high persistence of index option market-makers' variance risk exposure found is in line with the evidence in Madhavan and Smidt (1993), who show that the inventory of NYSE specialists can deviate from their long-run mean for several weeks.

As apparent in the bottom Panel of Figure 2 inventory risk is negative most of the time. Structurally, this stylized fact is captured by $\hat{\alpha}$ estimated to be large and negative. Note however, that the unconditional expectation of inventory risk is also influenced by $\hat{\psi}$. Interestingly, the loading of inventory risk on lagged spot variance is positive. Thus, inventory risk tends to increase with market uncertainty. This result is consistent with time series regressions of daily changes in inventory risk on the VIX index which also yield to a positive factor loading for the VIX index.

In-sample, the volatility parameter for inventory risk is 16.55%, which is of similar order of stock volatility. Note the way the instantaneous correlation of inventory risk with market variance is almost zero. While changes in inventory risk increases conditionally with market variance through ψ, inventory risk is nearly contemporaneously independent of market variance innovations.

In the model w captures the dealer's initial wealth. The estimated wealth level is about 440 million dollars. This parameter estimate is comparable in magnitude to the daily delta-hedged profits and losses documented in Figure 3. Quantitatively, it represents approximately 4 years of cumulative daily profit and losses.

In-Sample Performance

Finally, the model's sum of squared errors is about 6.68 at the optimum. To benchmark the model performance, we fit Heston (1993) on option prices using a similar methodology but for which we set the instantaneous $VRP_t = h \cdot V_t$, where h is a constant to be estimated. This specification for the variance risk premium is consistent with Heston (1993), Bates (2000), and Pan (2002) among others. For ease of comparison with the Heston model, IRW model (i.e. inventory risk and wealth model) will henceforth refer to our baseline model which dynamics satisfy Proposition 3.

From Panel B, we see that the variance risk premium parameter for the Heston model is -1.08. Note however that the fit obtained by the Heston model is not as good as the fit of the IRW model. During our sample, Heston's sum of squared implied volatility errors is about 20% greater than the IRW model.

5.3 Index Option Fit

To measure the performance of the two models, we compute the percentage implied volatility RMSE defined as

$$IVRMSE \equiv \sqrt{\frac{1}{N_J}\sum_{j,t}^{N_J}\left(IV_{j,t} - IV_{j,t}^M\right)^2} \times 100, \tag{5.4}$$

where $IV_{j,t}^M$ is the model implied volatility. We compute the performance criterion for both models using all puts options with moneyness between 0.9 and 1.1, and with 2, 3, and 6 months to maturity. The sample consists of $131,838$ observations, which is considerably larger than the sample used for the estimation.

Table 11 presents the fit obtained for the IRW model and Heston. We report the IVRMSE for all contracts averaged by year in Panel A, by moneyness in Panel B, and by maturity in Panel C. The average yearly IVRMSE for the IRW model is 3.14%, which is good given that the sample includes the financial crisis.

The fit of each model are close; however, the IRW model dominates the Heston model particularly in the second part of the sample period. Relative to the inventory model, Heston's best pricing performance is obtained in the period 1997-1999 with a IVRMSE difference between the two model of 0.43%. In contrast, the IRW model achieves its best fit relative to the Heston model in 2009-2011 with an IVRMSE lower by 1%. Since 2003, the IVRMSE of the inventory model is lower than Heston by more than 0.91%.

From Panel B in Table 11 we see that the inventory model achieves a better fit for ATM puts relative to the Heston model. ATM options are the most sensitive to changes in variance risk premium. Thus, this result potentially uncovers the inability of the Heston model combined with the specification $VRP_t = h \cdot V_t$ to adequately captures the variations in the variance risk premium. To further assess this, we regress the daily implied volatility root mean squared errors against the empirical proxy of one-month variance risk premium, and the daily log-likelihood $\mathcal{L}_t$ of the physical returns. For the IRW model, we obtain

$$IVRMSE_t = \underset{(0.00)}{3.14} + \underset{(0.90)}{0.02} \cdot VRP_{t,30} - \underset{(0.01)}{0.27} \cdot \mathcal{L}_t + \varepsilon_t, \tag{5.5}$$

where the regressors are standardized, and the Newey-West p-values reported under each parameters are calculated with 8 lags. The adjusted R-Squared of the regression is 1.77%. Arguably, the mispricing of the IRW model is not related to the realization of the variance risk premium. We see however that when the likelihood decreases the pricing error of the model tends to increase. For Heston (1993), we have

$$IVRMSE_t = \underset{(0.00)}{3.56} - \underset{(0.00)}{0.81} \cdot VRP_{t,30} - \underset{(0.75)}{0.02} \cdot \mathcal{L}_t + \varepsilon_t, \tag{5.6}$$

with an adjusted R-Squared of 17.89%. The result is striking. In contrast to the IRW model, a large part of the pricing errors of the Heston model can be attributed by its inability to fully capture the fluctuations in the variance risk premium.

We also see from Panel B in Table 11 that the IRW model fits OTM puts better than Heston. This result is partly due to the feedback effects between the variance risk premium, and the ratio of inventory risk to wealth. When the variance risk premium is negative, the ratio of inventory risk to wealth tends to decrease when volatility increases. This further reduces the variance risk premium and increases the level of risk-neutral variance during bad times. Because a high risk-neutral volatility during low market returns results in a negative skewness of the index returns distribution, this feedback effects help improve the pricing of OTM puts.

In Panel C of the same Table, we see that the IRW model achieves the best performance across maturity. The term structure of risk-neutral volatility is directly influenced by the term structure of variance risk premia. Thus, this finding is consistent with the results in Table 4. Because inventory risk and market-maker's wealth matter for the term structure of variance risk premia, they are also important for the term structure of risk-neutral volatility.

In Figure 5, we plot the daily one-month variance risk premium and IVRMSE for each model. From Panel A, we see that the IRW model can deliver a wide range of variance risk premium. In comparison, the one-month variance risk premium implied by the Heston model is always negative. The ability of the inventory model to generate substantial variation in the variance risk premium drives its performance.

Overall, the previous results suggest that the IRW model performs better than Heston (1993) within our sample. It is well known that models with more parameters often obtain better in-sample fit but lead to poorer out-of-sample performance. This may influence our finding since the IRW model has more parameters than Heston (1993). To assess this, we forecast on each day next's day latent variables fixing the structural parameters to their optimum and taking previous day latent variables as given. For the IRW model, we forecast spot inventory and spot variance on each day based on the dynamic (3.7) and (3.2). Similarly, we forecast the spot variance for Heston (1993) on each day taking the previous day filtered spot as given. Based on these predicted values and on the estimated parameters we assess the fit of each model. This constitutes a good sanity check for the inventory dynamic wich has been estimated purely from option prices. We present the results in Table 12. Despite the forecast errors, the IRW model still obtain a better performance than Heston (1993).

The results suggest that accounting for inventory risk and market-maker wealth is important to capture variation in the variance risk premium and thus explain index option prices.

In the next section, we quantify the economic magnitude of the response of SPX option prices to changes in inventory risk and market-maker's wealth.

5.4 Economic Magnitude

Little is known about the way market-makers adjust option mid-quotes when they absorb large buy orders, which cause their exposure to market variance to become more negative (i.e. decrease their inventory risk). Similarly, the magnitude of the impact of market-makers' losses and gains on index option prices is also an open question.

To understand the way inventory risk and market-maker's wealth influence market-makers' quotation behaviour, in Figure 6 we plot the dollar sensitivity of SPX put options to a decrease in the state variables implied by the model. To calculate the model sensitivies $\partial P_t/\partial InvRisk_t$ and $\partial P_t/\partial W_t$, we use the estimated parameters $\hat{\Theta}^V$, $\hat{\Theta}^{Inv}$, and $\hat{w}$, and set $r = 4\%$, $q = 0$, $S_t = 1183$, and $V_t = \hat{\theta}$. Moreover, inventory risk is set to its unconditional mean $InvRisk_t = -9.03E + 09$, and market-maker wealth is initialized to $W_t = \hat{w}$. Based on the sensitivities calculated, we then compute the dollar response of each option as $\Delta InvRisk_t \cdot \partial P_t/\partial InvRisk_t$ and $\Delta W_t \cdot \partial P_t/\partial W_t$, and plot the result across moneyness. In both panels, each line corresponds to different level of decrease in inventory risk and wealth. The circles correspond to the dollar response to an average decrease in the state variables. The diamonds identify the response when the decrease of each latent variable is within its 90^{th} percentile decrease.

Figure 6 provides several insights. First, we see that a decrease in inventory risk results in an increase in index option prices. This is consistent with the theoretical predictions in Proposition 2. When market-makers' risk exposure decreases, they require a more negative variance risk premium which translates into an increase in index option prices. Similarly, market-makers' losses result in an increase in the price of index options. Interestingly, market-maker's wealth has the biggest impact on option prices.

Note also that the effect of inventory risk and market-maker's wealth on SPX options across strike prices is non-linear, and is most prominent for at-the-money options. Interestingly, it is these options that make up most of market-makers' inventory (see Table 2).

When all variables are set to their mean values, the average decrease in inventory risk and the average wealth loss cause between a 10 and 50 dollars increase in price. Given an average option price of $6,500$ dollars, these effects correspond to a 0.15% and 0.77% daily increase in price.

During turbulent times, when market-makers' aggregate loss is within its 90^{th} percentile decrease, it causes a 150 dollars increase in price, which corresponds to a 2.31% daily increase. These results further highlight the non-trivial role of market-makers in the determination of index option prices through their effect on the variance risk premium.

6 Summary and Conclusions

Using aggregate market-makers' positions at the CBOE, we examine how inventory risk and market-makers' wealth jointly determine the value of index options through their effects on the variance risk premium. The analysis demonstrates that part of the variance risk premium corresponds to option market-makers' compensation for their exposure to market variance. We find that inventory exposure to market variance and changes in market-makers' wealth explain the variance risk premium. Using daily observations, we find that a one standard deviation decrease in inventory risk causes a 1.2% decrease in the variance risk premium next period. Moreover, the interaction of lagged inventory risk with changes in market-maker wealth is also strongly significant.

Based on these findings, we then develop a model in which market variance is stochastic and a representative market-maker with limited capital accumulates inventory over time by absorbing end-users' net demand for index options. Starting from the market-maker's optimal trading strategy, we derive an explicit formula linking the variance risk premium to inventory risk and market-maker wealth. The model predictions provides several insights that help understand empirical stylized facts.

Finally, we estimate the model on S&P 500 returns and options. Overall, the model performs well. Our results suggest that accounting for inventory risk and market-maker's wealth in pricing models can help for index option valuation. Based on the estimated parameters, we find that the effect of inventory risk and market-maker's wealth on SPX options across strike prices is non-linear, and is most prominent for at-the-money options.

Several issues are left for future research. First, it would be interesting to develop and test the implications of various inventory risk dynamics for pricing. Second, it would be also interesting to extend the model by allowing for jumps in the market price. Third, it would be particularly compelling to estimate the model implied market-maker's wealth daily and study its predictive abilities for equity options. Finally, investigating the pricing implications inventory risk and market-makers' wealth have for other derivative markets such as the CDS market would be also of significant interest.

Appendix

This appendix starts by presenting the strategy used to forecast for integrated physical variance. It then collects the proofs of the propositions and discusses the algorithms used for estimating the model.

A. Forecasting Expected Physical Variance

Suppose we want to estimate the T-days expected integrated physical variance on date t_0 that is $RV_{t_0,T}$. Using a rolling-window of 252 days with $t \in \{t_0 - 252, ..., t_0 - 1\}$, we run the following HAR-RV model in the spirit of Corsi (2009)

$$\begin{aligned} \ln(RV_{t,T}) &= a_0 + a_1 \ln(RV_{t-1,1}) + a_2 \ln(RV_{t-6,6}) + \\ &\quad a_3 \ln(RV_{t-30,30}) + a_4 \ln(RV_{t-60,60}) + a_5 \ln(RV_{t-90,90}) \\ &\quad + a_6 \ln(RV_{t-120,120}) + \varepsilon_{t,30}, \end{aligned} \tag{6.1}$$

where $RV_{t,T}$ is the T-days realized variance at time t. We can write the model in matrix form. We have $\ln(RV_{t,T}) = X_{t-1} \cdot A + \varepsilon_{t,T}$ where X_{t-1} contains the explanatory variables known on day $t-1$ and A is the matrix of parameters. Using the OLS estimate for $\hat{A}$ and setting the regressors to their values on the day t_0, the model prediction for $RV_{t_0,T}$ on that day is $\exp(X_{t_0} \cdot \hat{A} + \frac{\hat{\sigma}^2_\varepsilon}{2})$ where $\hat{\sigma}^2_\varepsilon$ is the variance of the residuals. We repeat this procedure on each day and for each horizon.

B. Proof of Proposition 1

For ease of notation, we define $a_t^P \equiv E_t^P[\frac{dV_t}{dt}]$ and $a_t^Q \equiv E_t^Q[\frac{dV_t}{dt}]$ the physical and risk-neutral market variance drifts respectively. Applying Ito's lemma to f^j implies the following dynamic of index options under the P-measure

$$df_t^j = \left(\frac{\partial f_t^j}{\partial t} + \frac{\partial f_t^j}{\partial S_t} S_t(\mu - q) + \frac{\partial f_t^j}{\partial V_t} a_t^P + \frac{\partial^2 f_t^j}{\partial S_t \partial V_t} \rho_V \delta S_t V_t^{\eta+0.5} + \frac{1}{2}\left(\frac{\partial^2 f_t^j}{(\partial S_t)^2} V_t S_t^2 + \frac{\partial^2 f_t^j}{(\partial V_t)^2} (\delta V_t^\eta)^2\right)\right) dt + \frac{\partial f_t^j}{\partial S_t} S_t \sqrt{V_t}\left(\sqrt{1-\rho_V^2} dZ_t^S + \rho_V dZ_t^V\right) + \frac{\partial f_t^j}{\partial V_t} \delta V_t^\eta dZ_t^V. \quad (6.2)$$

We also know that given the dynamics (3.1) to (3.4), the price of any derivative f^j must satisfy the PDE

$$r f_t^j = \frac{\partial f_t^j}{\partial t} + \frac{\partial f_t^j}{\partial S_t} S_t(r - q) + \frac{\partial f_t^j}{\partial V_t} a_t^Q + \frac{\partial^2 f_t^j}{\partial S_t \partial V_t} \rho_V \delta S_t V_t^{\eta+0.5} + \frac{1}{2}\left(\frac{\partial^2 f_t^j}{(\partial S_t)^2} V_t S_t^2 + \frac{\partial^2 f_t^j}{(\partial V_t)^2} (\delta V_t^\eta)^2\right). \quad (6.3)$$

Combining (6.2) with (6.3), we obtain

$$df_t^j = \left(r f_t^j + \frac{\partial f_t^j}{\partial S_t} S_t(\mu - r - q) + \frac{\partial f_t^j}{\partial V_t}\left(a_t^P - a_t^Q\right)\right) dt + \frac{\partial f_t^j}{\partial S_t} S_t \sqrt{V_t}\left(\sqrt{1-\rho_V^2} dZ_t^S + \rho_V dZ_t^V\right) + \frac{\partial f_t^j}{\partial V_t} \delta V_t^\eta dZ_t^V. \quad (6.4)$$

Since the delta replication of f^j satisfies $\Delta Rep_t^j = \theta_t^B \cdot B_t + \theta_t^S \cdot S_t$ where θ^B and θ^S denote the units of bond and market index to hold. For an investment horizon dt, we have $\theta_t^S \cdot S_t = \frac{\partial f_t^j}{\partial S_t} S_t$ and $\theta_t^B \cdot B_t = f_t^j - \frac{\partial f_t^j}{\partial S_t} S_t$. Consequently, the replication portfolio when dividends are reinvested evolves as

$$\begin{aligned} d\Delta Rep_t^j &= \theta_t^B \cdot dB_t + \theta_t^S \cdot (dS_t + q S_t dt) \\ &= \left(f_t^j - \frac{\partial f_t^j}{\partial S_t} S_t\right) \cdot \frac{dB_t}{B_t} + \left(\frac{\partial f_t^j}{\partial S_t} S_t\right) \cdot \left(\frac{dS_t}{S_t} + q dt\right), \end{aligned} \quad (6.5)$$

with $\Delta Rep_t^j = f_t^j$. Combining (6.4) and (6.5), we get

$$\begin{aligned} df_t^j &= d\Delta Rep_t^j + \vartheta_t^j \cdot \left(VRP_t dt + \delta V_t^\eta dZ_t^V\right) \\ &= d\Delta Rep_t^j + \vartheta_t^j \cdot dF_t^V, \end{aligned} \quad (6.6)$$

where $VRP_t \equiv a_t^P - a_t^Q = \frac{1}{dt}\left(E_t^P[V_{t+dt}] - E_t^Q[V_{t+dt}]\right) = \delta V_t^\eta \cdot \phi_t^V$, $\vartheta_t^j \equiv \frac{\partial f_t^j}{\partial V_t}$, and $dF_t^V = VRP_t dt + \delta V_t^\eta dZ_t^V$ is the market variance risk factor. We can now express ϑ_t^j in terms of sensitivity to volatility

$$\vartheta_t^j = \frac{\partial f_t^j}{\partial V_t} = \frac{\partial f_t^j}{\partial \sqrt{V_t}} \cdot \frac{\partial \sqrt{V_t}}{\partial V_t} = Vega_t^j \cdot \frac{1}{2\sqrt{V_t}}, \quad (6.7)$$

which completes the proof.

C. Proof of Proposition 2

We adopt the following strategy. First, we solve the market-maker's (unconstrained) portfolio allocation when the market clearing condition is not imposed. Based on the investment strategy obtained, we then require it to satisfy the market clearing condition and infer the structure of the variance risk premium.

The static maximization

$$\max_{\tilde{\pi}_t} \; E^P[\ln(W_T)] \quad \text{subject to} \quad \begin{array}{l} \cdot \; W_0 \geqslant E^P[\xi_T W_T] \\ \cdot \; W_t \geqslant 0, \end{array} \tag{6.8}$$

is the dual problem of the unconstrained utility maximization (3.9) (see, among others, Karatzas, Lehoczky, and Shreve, 1987, and Cox, and Huang, 1989). To solve this, we form the lagrangian

$$\begin{aligned} L(\gamma) &= E^P[\ln(W_T)] + \gamma(W_0 - E^P[\xi_T W_T]) \\ &= E^P[\ln(W_T) - \gamma \xi_T W_T] + \gamma W_0, \end{aligned} \tag{6.9}$$

where γ is the lagrangian coefficient of the static budget constraint $W_0 = E^P[\xi_T W_T]$. A point wise maximization of (6.9) implies the following FOC condition

$$\frac{1}{W_t} = \gamma \xi_t. \tag{6.10}$$

Since the previous equation is valid for any t and $\xi_0 = 1$, the lagrangian coefficient satisfies $\gamma = 1/W_0$. Thus, optimally we have $W_t = \frac{W_0}{\xi_t}$. Given the definition of $\pi_t^{f,j}$, we also have $\pi_t^{f,j} W_t = \Phi_t^{MM,j} f_t^j \Longleftrightarrow \pi_t^{f,j} = \Phi_t^{MM,j} f_t^j / W_t$. Using this with (6.5) and (6.6) in Appendix B, we can write the market-maker's aggregate position in index options as

$$\begin{aligned} & \sum_j \pi_t^{f,j} \cdot \frac{df_t^j}{f_t^j} \\ = \; & \sum_j \frac{\Phi_t^{MM,j}}{W_t} \cdot df_t^j \\ = \; & \frac{dB_t}{B_t} \cdot \sum_j \frac{\Phi_t^{MM,j}}{W_t} \cdot \left(f_t^j - \frac{\partial f_t^j}{\partial S_t} S_t \right) + \left(\frac{dS_t}{S_t} + q dt \right) \cdot \sum_j \frac{\Phi_t^{MM,j}}{W_t} \cdot \left(\frac{\partial f_t^j}{\partial S_t} S_t \right) + \frac{dF_t^V}{2\sqrt{V_t}} \cdot \frac{InvRisk_t}{W_t}, \end{aligned} \tag{6.11}$$

where we use the definition (2.5) to uncover the market-maker's inventory risk. We can now use previous result to express the wealth process (3.8) as

$$\begin{aligned}\frac{dW_t}{W_t} &= \pi_t^B \cdot \frac{dB_t}{B_t} + \pi_t^S \cdot \left(\frac{dS_t}{S_t} + qdt\right) + \sum_j \pi_t^{f,j} \cdot \frac{df_t^j}{f_t^j} \\ &= \bar{\pi}_t^B \cdot \frac{dB_t}{B_t} + \bar{\pi}_t^S \cdot \left(\frac{dS_t}{S_t} + qdt\right) + \frac{InvRisk_t}{W_t} \cdot \frac{dF_t^V}{2\sqrt{V_t}},\end{aligned} \tag{6.12}$$

where $\bar{\pi}_t^B = \pi_t^B + \sum_j \frac{\Phi_t^{MM,j}}{W_t}\left(f_t^j - \frac{\partial f_t^j}{\partial S_t}S_t\right)$ and $\bar{\pi}_t^S = \pi_t^S + \sum_j \frac{\Phi_t^{MM,j}}{W_t}\left(\frac{\partial f_t^j}{\partial S_t}S_t\right)$ represent market-maker's total investment in the bond and in the index respectively.

In this economy, the discounted wealth process satisfies

$$\frac{d(\xi_t W_t)}{\xi_t W_t} = \frac{dW_t}{W_t} + \frac{d\xi_t}{\xi_t} + \frac{d\langle \xi, W\rangle_t}{\xi_t W_t}, \tag{6.13}$$

where $d\langle ., .\rangle_t$ is the covariance operator. Applying the previous equation on the SDF dynamic (3.4) and the wealth process (6.12), we obtain

$$\frac{d(\xi_t W_t)}{\xi_t W_t} = \left(\bar{\pi}_t^S \sqrt{V_t}\sqrt{1-\rho_V^2} - \phi_t^S\right) dZ_t^S + \left(\bar{\pi}_t^S \sqrt{V_t}\rho_V + 0.5\delta V_t^{\eta-0.5}\frac{InvRisk_t}{W_t} - \phi_t^V\right) dZ_t^V, \tag{6.14}$$

for which we have imposed the martingale conditions $\xi_t S_t + \int_0^t q\xi_u S_u du = E_t^P[\xi_T S_T + \int_0^T q\xi_u S_u du]$ and $\xi_t f_t^j = E_t^P[\xi_T f_T^j]$. We can now integrate (6.14) to express $\xi_T W_T$ in its integral form

$$\xi_T W_T = W_0 + \int_0^T \xi_t W_t \left(\left(\bar{\pi}_t^S \sqrt{V_t}\sqrt{1-\rho_V^2} - \phi_t^S\right) dZ_t^S + \left(\bar{\pi}_t^S \sqrt{V_t}\rho_V + 0.5\delta V_t^{\eta-0.5}\frac{InvRisk_t}{W_t} - \phi_t^V\right) dZ_t^V\right). \tag{6.15}$$

By application of the Clark-Ocone formula to $\xi_t W_t$, we also have the following expression for $\xi_T W_T$ in terms of its Malliavin derivatives

$$\begin{aligned}\xi_T W_T &= E_0^P[\xi_T W_T] + \int_0^T E_t^P[D_t^S(\xi_T W_T)]dZ_t^S + \int_0^T E_t^P[D_t^V(\xi_T W_T)]dZ_t^V \\ &= W_0 + \int_0^T E_t^P[D_t^S(\xi_T W_T)]dZ_t^S + \int_0^T E_t^P[D_t^V(\xi_T W_T)]dZ_t^V,\end{aligned} \tag{6.16}$$

where $D_t^i(X)$ is the time t Malliavin derivative of X with respect to Z^i for $i \in \{S, V\}$.[15] This repre-

[15] Malliavin derivatives have also been used to obtain explicit formulas for optimal investment strategies in Detemple, Garcia, and Rinsdisbacher (2003) and Detemple and Rinsdisbacher (2010) among others. We refer to the Appendix D in Detemple, Garcia, and Rindisbacher (2003) for an introduction to Malliavin calculus and a presentation of the Clarck-Ocone formula. We refer to Di Nunno, Økskendal, and Proske (2009) for an extensive treatment of Malliavin Calculus applied to Finance.

sentation of $\xi_T W_T$ can be combined with (6.15) to obtain explicit formulas for $\bar{\pi}^S$ and $InvRisk_t/W_t$. Because both (6.16) and (6.15) uniquely defined $\xi_T W_T$ the integrands in both equations must be equal. This implies

$$\bar{\pi}_t^S \sqrt{V_t}\sqrt{1-\rho_V^2} - \phi_t^S = E_t^P[D_t^S(\xi_T W_T)] \tag{6.17}$$

$$\bar{\pi}_t^S \sqrt{V_t}\rho_V + 0.5\delta V_t^{\eta-0.5}\left(\frac{InvRisk_t}{W_t}\right) - \phi_t^V = E_t^P[D_t^V(\xi_T W_T)]. \tag{6.18}$$

Together, the two previous equations define the market-maker's optimal investment strategy. We can now impose the market clearing condition to (6.17) and (6.18). The market clearing condition imposes $\Phi_t^{MM,j} = -\Phi_t^{EU,j}$ for all j and thus $InvRisk_t = -\sum_j \Phi_t^{EU,j} Vega_t^j$ in aggregate. Solving for $\bar{\pi}^S$ in (6.17) and using the result in (6.18), we get

$$\left(E_t^P[D_t^S(\xi_T W_T)] + \phi_t^V\right) - m\left(E_t^P[D_t^V(\xi_T W_T)] + \phi_t^S\right) = 0.5\delta V_t^{\eta-0.5}\left(\frac{InvRisk_t}{W_t}\right), \tag{6.19}$$

where $m \equiv \rho_V/\sqrt{1-\rho_V^2}$.[16] By the properties of Malliavin derivatives, we have for $i \in \{S, V\}$

$$D_t^i(\xi_T \cdot W_T) = D_t^i(\xi_T \cdot \frac{W_0}{\xi_T}) = D_t^i(W_0) = 0, \tag{6.20}$$

where we used the optimality condition $W_T = W_0/\xi_T$, and the fact that the Malliavin derivative of an adapted process is 0 (i.e. $D_t(X_s) = 0$ when $s < t$). Therefore, $E_t^P[D_t^S(\xi_T W_T)] = E_t^P[D_t^V(\xi_T W_T)] = 0$. Using this into (6.19), we see that

$$\phi_t^V - m\phi_t^S = 0.5\delta V_t^{\eta-0.5}\left(\frac{InvRisk_t}{W_t}\right). \tag{6.21}$$

relates market-maker's optimal inventory risk to the two prices of risks. The market index no-arbitrage imposes

$$\xi_t S_t + \int_0^t q\xi_u S_u du = E_t^P[\xi_T S_T + \int_0^T q\xi_u S_u du] \Leftrightarrow \frac{\mu - r}{\sqrt{V_t}} = \sqrt{1-\rho_V^2}\cdot\phi_t^S + \rho_V \cdot \phi_t^V. \tag{6.22}$$

We can use previous equation in order to express the market price of risks ϕ^S in terms of the market premium and ϕ^V

$$\phi_t^S = \frac{\mu - r}{\sqrt{V_t(1-\rho_V^2)}} - m\cdot\phi_t^V. \tag{6.23}$$

[16]Note that m is well-defined whenever $abs(\rho_V) \neq 1$.

Combining (6.21) and (6.23), we can express the market price of variance risk as

$$\phi_t^V = \rho_V \left(\frac{\mu - r}{\sqrt{V_t}}\right) + 0.5(1-\rho_V^2)\delta V_t^{\eta-0.5}\left(\frac{InvRisk_t}{W_t}\right). \tag{6.24}$$

Given that the variance risk premium satisfies $VRP_t = (\delta V_t^\eta)\cdot \phi_t^V$, we finally get

$$VRP_t = \rho_V \delta V_t^\eta \left(Sharpe_t\right) + 0.5(1-\rho_V^2)\delta^2 V_t^{2\eta-0.5}\left(\frac{InvRisk_t}{W_t}\right), \tag{6.25}$$

where $Sharpe_t = \left(\frac{\mu - r}{\sqrt{V_t}}\right)$, and $InvRisk_t = -\sum_j \Phi_t^{EU,j} Vega_t^j$.

D. Particle Filter Estimation

The following algorithm describes the way we evaluate the likelihood $\tilde{P}_t^j\left(V_t^j, \Theta^V\right)$ of observing S_{t+1} given the smooth resampled particles V_t^j, and the structural parameters Θ^V. For estimation purposes, we set the number of particles denoted N to $10,000$.

Using the Euler discretization for $dln(S_t)$ and (3.2), one can simulate the state of the N raw particles $\left\{\tilde{V}_t^j\right\}_{j=1}^N$ forward given $\left\{V_{t-1}^j\right\}_{j=1}^N$ according to

$$Z_t^{V,j} = \left(\frac{\ln\left(\frac{S_t}{S_{t-1}}\right) - \left(\mu - q - \frac{V_t^j}{2}\right)\Delta t}{\sqrt{V_t^j}} - \sqrt{1-\rho_V^2}\, Z_t^{S,j}\right)/\rho_V \tag{6.26}$$

$$\tilde{V}_t^j = V_{t-1}^j + \kappa(\theta - V_{t-1}^j)\Delta t + \delta\left(V_{t-1}^j\right)^\eta Z_t^{V,j}, \tag{6.27}$$

where $Z_t^{S,j}$ is $N(0,\sqrt{\Delta t})$, and μ is fixed to the sample average. Using the set of raw particles, the likelihood of observing S_{t+1} given $\tilde{V}_t^j$ and S_t is

$$\tilde{P}_t^j\left(\tilde{V}_t^j, \Theta^V\right) = \frac{1}{\sqrt{2\pi\tilde{V}_t^j}}\exp\left(-\frac{\left(\ln\left(\frac{S_{t+1}}{S_t}\right) - \left(\mu - q - \frac{\tilde{V}_t^j}{2}\right)\Delta t\right)^2}{2\tilde{V}_t^j}\right). \tag{6.28}$$

Based on the set of normalized weights

$$\breve{P}_t^j\left(\tilde{V}_t^j, \Theta^V\right) = \frac{\tilde{P}_t^j\left(\tilde{V}_t^j, \Theta^V\right)}{\sum\limits_j \tilde{P}_t^j\left(\tilde{V}_t^j, \Theta^V\right)}, \tag{6.29}$$

and the raw $\tilde{V}_t^j$, the method of Pitt (2002) can be applied to resample the smoothed particles $\left\{V_t^j\right\}_{j=1}^N$ and evaluate their corresponding weights $\tilde{P}_t^j\left(V_t^j, \Theta^V\right)$.[17]

E. Risk-Neutral Pricing

Suppose that we want to price an index put option on day t with T days to maturity and strike price K based on $N = 10,000$ simulations. For each simulation n, we initiate the state variables S_t, V_t, and $InvRisk_t$ to their respective values on the day of the pricing. Moreover, we initialize the market-maker wealth to w. For a given path n, the forward state of the discretized processes in Proposition 3 given the information on day t is

$$\ln(S_{t+1}^n) = \ln(S_t^n) + (r - q - V_t^n/2)\,\Delta t + \sqrt{V_t^n}\left(\sqrt{1-\rho_V^2}\tilde{Z}_{t+1}^{S,n} + \rho_V \tilde{Z}_{t+1}^{V,n}\right) \tag{6.30}$$

$$V_{t+1}^n = V_t^n + \kappa(\theta - V_t^n)\Delta t - VRP_t^n \Delta t + \delta\,(V_t^n)^\eta\,\tilde{Z}_{t+1}^{V,n} \tag{6.31}$$

$$\begin{aligned} InvRisk_{t+1}^n \;=\;& InvRisk_t^n + \lambda(\alpha - InvRisk_t^n)\Delta t + \psi V_t^n \Delta t + \sigma InvRisk_t^n\left(\sqrt{1-\rho_{Inv}^2}\tilde{Z}_{t+1}^{S,n} + \rho_{Inv}\tilde{Z}_{t+1}^{V,n}\right) \\ & -\sigma InvRisk_t^n\left(\sqrt{1-\rho_{Inv}^2}\phi_t^{S,n} + \rho_{Inv}\phi_t^{V,n}\right)\Delta t \end{aligned} \tag{6.32}$$

$$\ln(W_{t+1}^n) = \ln(W_t^n) + \left(r - \left(\left(\phi_t^{S,n}\right)^2 + \left(\phi_t^{V,n}\right)^2\right)/2\right)\Delta t + \phi_t^{S,n}\tilde{Z}_{t+1}^{S,n} + \phi_t^{V,n}\tilde{Z}_{t+1}^{V,n}, \tag{6.33}$$

where $\tilde{Z}_{t+1}^{S,n}$ and $\tilde{Z}_{t+1}^{V,n}$ are independent $N(0,\sqrt{\Delta t})$. In the previous system, we set

$$VRP_t^n = \rho_V \delta\,(V_t^n)^\eta\,Sharpe_t^n + 0.5(1-\rho_V^2)\delta^2\,(V_t^n)^{2\eta-0.5}\left(\frac{InvRisk_t^n}{W_t^n}\right), \tag{6.34}$$

where $Sharpe_t^n = \left(\frac{\mu - r}{\sqrt{V_t^n}}\right)$. Moreover, the prices of risks are calculated according to

$$\phi_t^{S,n} = Sharpe_t^n/\sqrt{1-\rho_V^2} - \rho_V\phi_t^{V,n}/\sqrt{1-\rho_V^2} \quad \text{and} \quad \phi_t^{V,n} = VRP_t^n/\left(\delta\,(V_t^n)^\eta\right). \tag{6.35}$$

[17]The method proposed in Pitt (2002) involves smoothing the $\breve{P}_t^j$ to a continuous CDF from which the set of smooth particle V_t^j can be resampled.

Simulating the system forward from day t to T, the price of the index put option on day t is equal to

$$P\left(\Theta^{Inv}, w, \Theta^{V}, V_t, InvRisk_t\right) = \sum_{n=1}^{N} \frac{\max\left(K - S_T^n, 0\right) \cdot \exp(-r \cdot (T-t)/365)}{N}, \tag{6.36}$$

where Θ^V and Θ^{Inv} are the structural parameters of the market variance and inventory risk processes respectively, and w is the market-maker's wealth parameter.

References

[1] Adrian, T., Etula, E., Muir, T. (2014) Financial Intermediaries and the Cross-section of Asset Returns, Journal of Finance, 69, 2557–2596.

[2] Adrian, T., Shin, H.S. (2010) The changing nature of financial intermediation and the financial crisis of 2007–2009, Annual Review of Economics, 2, 603–618.

[3] Aït-Sahalia, Y., Karaman, M., Mancini, L. (2013) The term structure of variance swaps and risk premia, Working Paper, Princeton University.

[4] Aït-Sahalia, Y., Kimmel, R. (2007) Maximum likelihood estimation of stochastic volatility models, Journal of Financial Economics, 83, 413–452.

[5] Amihud, Y., Mendelson, H., (1980) Dealership market: Market making with inventory, Journal of Financial Economics, 8, 31–53.

[6] Ang, A., Hodrick, R., Xing, Y., Zhang, X. (2006) The cross-section of volatility and expected returns, Journal of Finance, 61, 259–299.

[7] Bakshi, G., Kapadia, N. (2003) Delta-hedged gains and the negative market volatility risk premium, Review of Financial Studies, 16, 527–566.

[8] Bakshi, G., Madan, D. (2006) A Theory of volatility spread, Management Science, 52, 1945–1956.

[9] Barras, L., Malkhozov, A. (2014) When variance risk has two prices: Evidence from the equity and option markets, Working Paper, McGill University.

[10] Bates, D. (2000) Post-'87 crash fears in the S&P 500 futures option market, Journal of Econometrics, 94, 181–238.

[11] Bates, D. (2003) Empirical option pricing: A retrospection, Journal of Econometrics, 116, 387–404.

[12] Black, F., Scholes, M. (1973) Valuation of options and corporate liabilities, Journal of Political Economy, 81, 637–654.

[13] Bollen, N.P., Whaley, R.E. (2004) Does net buying pressure affect the shape of implied volatility functions?, Journal of Finance, 59, 711–53.

[14] Bollerslev, T., Tauchen, G., Zhou, H. (2009) Expected stock returns and variance risk premia, Review of Financial Studies, 22, 4463–4492.

[15] Britten-Jones, M., Neuberger, A. (2000) Option prices, implied price processes, and stochastic volatility, Journal of Finance, 55, 839–866.

[16] Brunnermeier, M., Pedersen, L. (2009) Market liquidity and funding liquidity, Review of Financial Studies, 22, 2201–2238.

[17] Buraschi, A., Trojani, F., Vedolin, A. (2014) When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia, Journal of Finance, Forthcoming.

[18] Carr, P., Wu, L. (2007) Stochastic skew in currency options, Journal of Financial Economics, 86, 213–247.

[19] Carr, P., Wu, L. (2009) Variance risk premiums, Review of Financial Studies, 22, 1311–1341.

[20] Chen, H., Joslin, S., Ni, S. (2013) Demand for crash insurance and stock returns, Working Paper, MIT.

[21] Christoffersen, P., Jacobs, K., Mimouni, K. (2010) Volatility dynamics for the S&P 500: evidence from realized volatility, daily returns, and option prices, Review of Financial Studies, 23, 3141–3189.

[22] Comerton-Forde, C., Hendershott, T., Jones, C., Moulton, P., Seasholes, M. (2010) Time variation in liquidity: The role of market-maker inventories and revenues, Journal of Finance 65, 295–331.

[23] Corsi, F. (2009) A simple approximate long-memory model of realized volatility, Journal of Financial Econometrics, 7, 174–196.

[24] Cox, J.C., Huang, C.F. (1989) Optimal consumption and portfolio plicies when asset prices follow a diffusion process, Journal of Economic Theory, 49, 33–83.

[25] Cremers, M., Halling, M., Weibaum, D. (2014) Aggregate jump and volatility risk in the cross-section of stock returns, Journal of Finance, Forthcoming.

[26] Detemple, J. B., Garcia, R., Rindisbacher, M. (2003) A Monte carlo method for optimal portfolios, Journal of Finance, 58, 401–446.

[27] Detemple, J. B., Garcia, R., Rindisbacher, M. (2005) Intertemporal asset allocation: a comparison of methods, Journal of Banking and Finance, 29, 2821–2848.

[28] Detemple, J. B., Rindisbacher, M. (2010) Dynamic asset allocation: portfolio decomposition formula and applications, Review of Financial Studies, 23, 25–100.

[29] Di Nunno, G., Øksendal, B., Proske, F. (2009) Malliavin calculus for Lévy processes with applications to finance, Springer-Verlag.

[30] Driessen, J., Maenhout, P., Vilkov, G. (2009) The price of correlation risk: Evidence from equity options, Journal of Finance, 64, 1377–1406.

[31] Egloff, D., Leippold, M., and Wu, L. (2010) The term structure of variance swap rates and optimal variance swap investments, Journal of Financial and Quantitative Analysis, 45, 1279–1310.

[32] Elkamhi, R., Stefanova, D. (2011) Dynamic correlation or tail dependence hedging for portfolio selection, Working Paper, University of Toronto.

[33] Gârleanu, N., Pedersen, L.H., Poteshman, A. (2009) Demand-based option pricing, Review of Financial Studies, 22, 4259–4299.

[34] Gromb, D., Vayanos, D. (2002) Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs, Journal of Financial Economics, 66, 361–407.

[35] Hansch, O., Naik, N., Viswanathan, S., (1998) Do inventories matter in dealership markets? Evidence from the London Stock Exchange, Journal of Finance, 53, 1623–1655.

[36] Heston, S. (1993) Closed-form solution for options with stochastic volatility, with application to bond and currency options, Review of Financial Studies, 6, 327–343.

[37] Ho, T., Stoll, H., (1981) Optimal dealer pricing under transactions and return uncertainty, Journal of Financial Economics, 9, 47–73.

[38] Ho, T., Stoll, H., (1983) The dynamics of dealer markets under competition, Journal of Finance, 38, 1053–1074.

[39] Johannes, M., Polson, N., Stroud, J. (2009) Optimal filtering of jump-diffusions: Extracting latent states from asset prices, Review of Financial Studies, 22, 2759–2799.

[40] Jones, C.S. (2003) The dynamics of stochastic volatility: evidence from underlying and options markets, Journal of Econometrics, 116, 181–224.

[41] Karatzas, I., Lehoczky, J.P., Shreve, S. (1987) Optimal portfolio and consumption decisions for a small investor on a finite time-horizon, Journal on Control and Optimization, 25, 1557–1586.

[42] Leippold, M., Su, L. (2011) Collateral smile, Working Paper, Swiss Finance Institute.

[43] Madhavan, A., Smidt, S., (1993) An analysis of daily changes in specialist inventories and quotations, Journal of Finance, 48, 1595–1628.

[44] Madhavan, A., Sofianos, G., (1998) An empirical analysis of NYSE specialist trading, Journal of Financial Economics, 48, 189–210.

[45] Malik, S., Pitt, M. (2011) Particle filters for continuous likelihood evaluation and maximisation, Journal of Econometrics, 165, 190–209.

[46] Mildenstein, E., Schleef, H., (1983) The optimal pricing policy of a monopolistic market maker in the equity market, Journal of Finance, 38, 218–231.

[47] Naik, N., Yadav, P., (2003) Do dealer firms manage inventory on a stock-by-stock or a portfolio basis? Journal of Financial Economics, 69, 325–353.

[48] Pan, J. (2002) The jump risk premia implicit in options: Evidence from an integrated time series study, Journal of Financial Economics, 63, 3–50.

[49] Pitt, M. (2002) Smooth particle filters for likelihood evaluation and maximization, Manuscript, University of Warwick.

[50] Shleifer, A., Vishny, R. (1997) The Limits of Arbitrage, Journal of Finance, 36, 35–55.

[51] Todorov, V. (2010) Variance risk premia dynamics: The role of jumps, Review of Financial Studies, 23, 345–383.

[52] Trolle, A., Schwartz, E. (2009) Unspanned stochastic volatility and the pricing of commodity derivatives, Review of Financial Studies, 22, 4423–4461.

[53] Vilkov, G. (2008) Variance risk premium demystified, Working Paper, INSEAD.

[54] Zhang, L., Mykland, P., Aït-Sahalia, Y. (2005) A Tale of two time scales: Determining integrated volatility with noisy high-frequency data, Journal of the American Statistical Association, 100, 1394–1411.

Figure 1: S&P 500 Index, Variance Risk Premium, and Delta-Hedged Gains and Losses of Near-the-Money Options

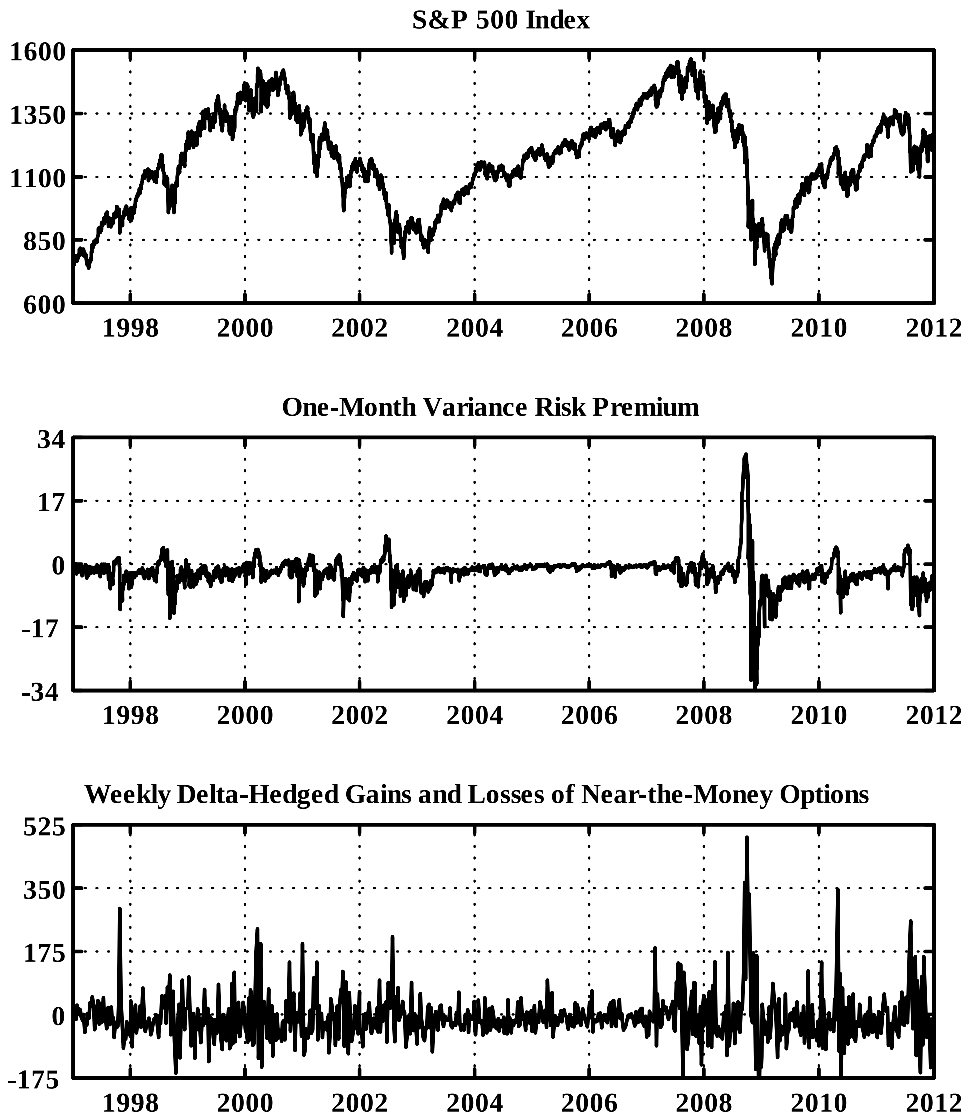

Notes to Figure: The top panel plots the time-series of the S&P 500 index. The middle panel plots the one-month variance risk premium expressed in percentage and measured by the difference between the one-month ex-post realized variance and one-month expected risk-neutral variance. In the bottom panel, we graph the weekly average of daily delta-hedged gains and losses of all options with moneyness (S/K) between 0.98 and 1.02.

Figure 2: CBOE VIX Index and
Market-Makers' Inventory Risk

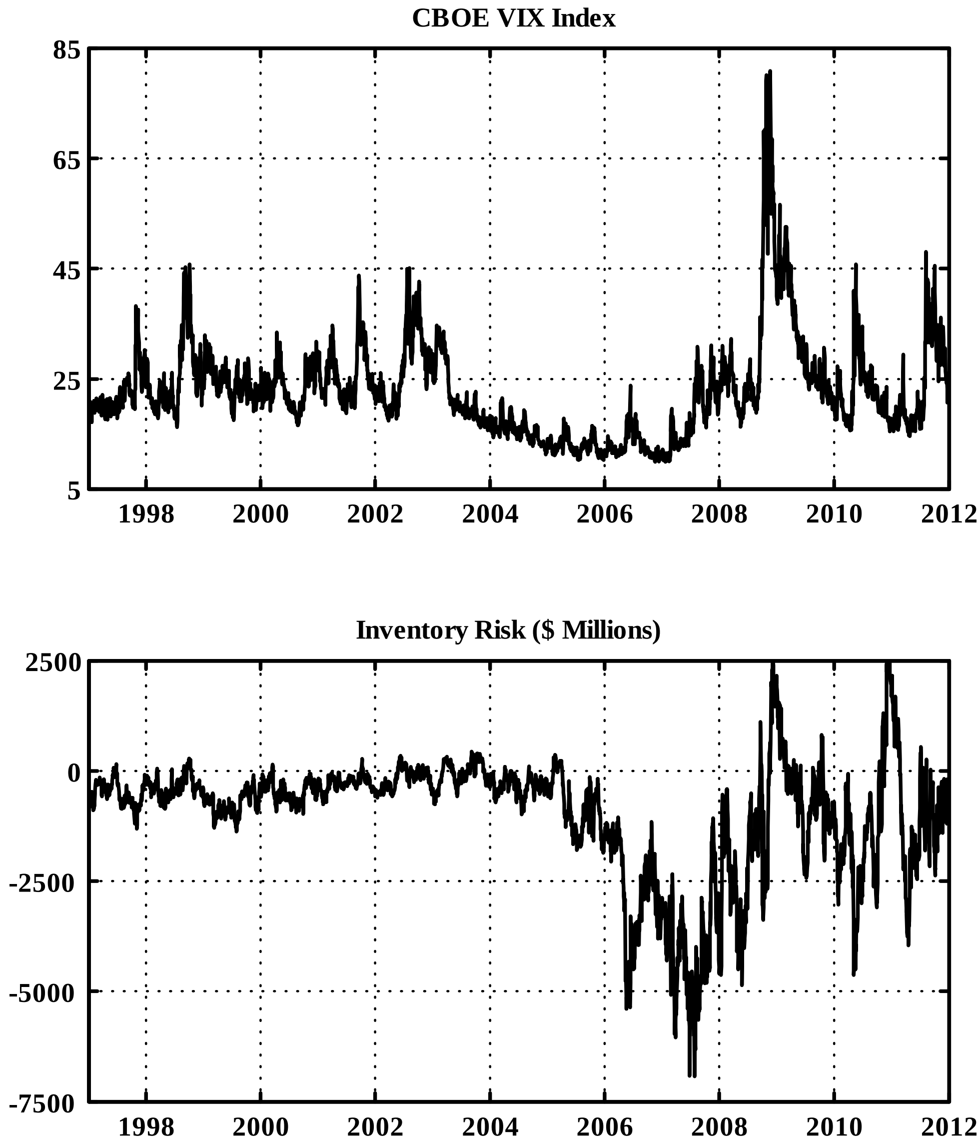

Notes to Figure: The top panel plots the CBOE VIX Index which represents the implied volatility of an at-the-money option with exactly 30 days to maturity expressed in percentage. The bottom panel plots CBOE market-makers' inventory risk dynamic measured as the vega-weighted sum of inventories across all contracts expressed in $ Millions.

Figure 3: Market-Makers' Daily and Cumulative Profits and Losses, and Market-Makers' Bid-Ask Spread Revenue

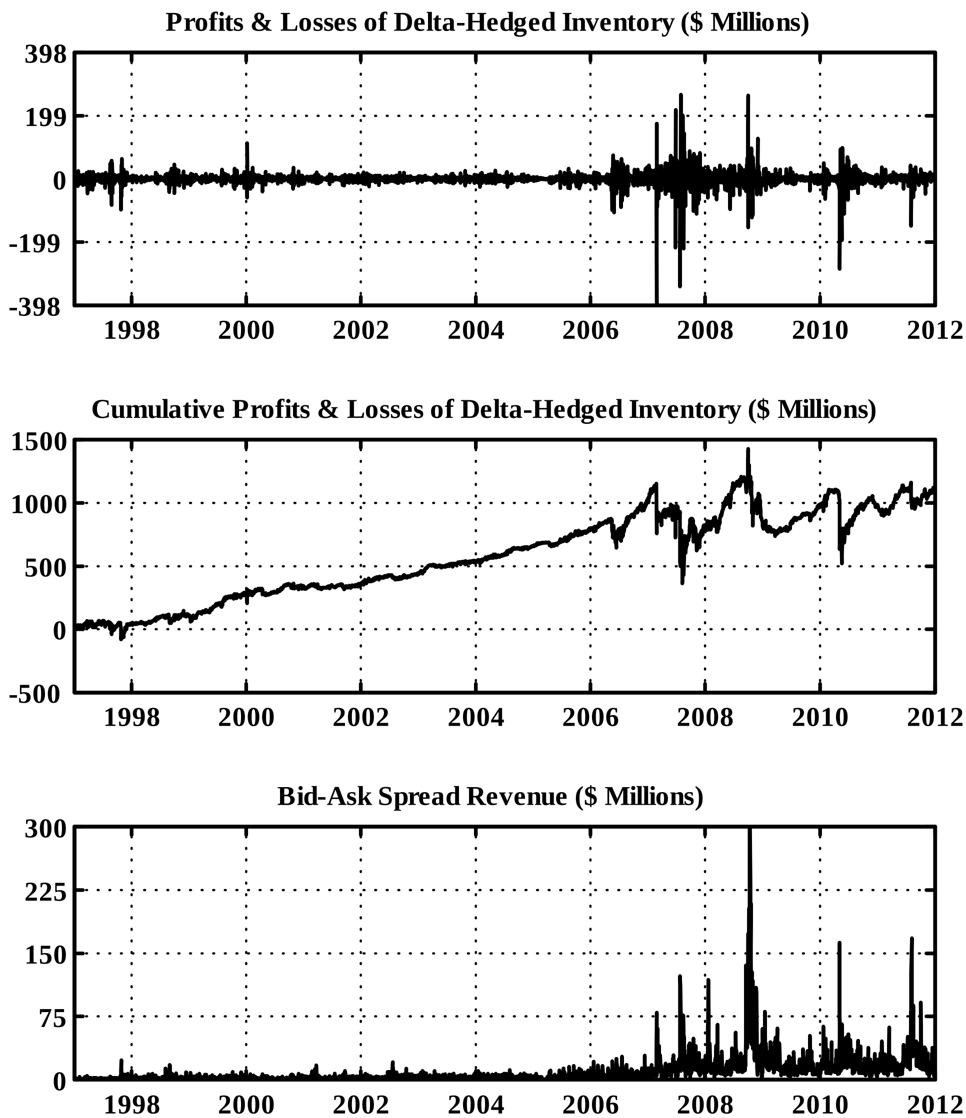

Notes to Figure: The top panel plots the daily profits and losses of market-makers' delta-hedged inventory expressed in $ Millions. The middle panel graphs the cumulative daily profits and losses over the sample period of market-makers' delta-hedged inventory in $ Millions. In the bottom panel, we plot the SPX market-makers' bid-ask spread revenue expressed in $ Millions.

Figure 4: Filtered Spot Volatilities Using Daily S&P 500 Returns

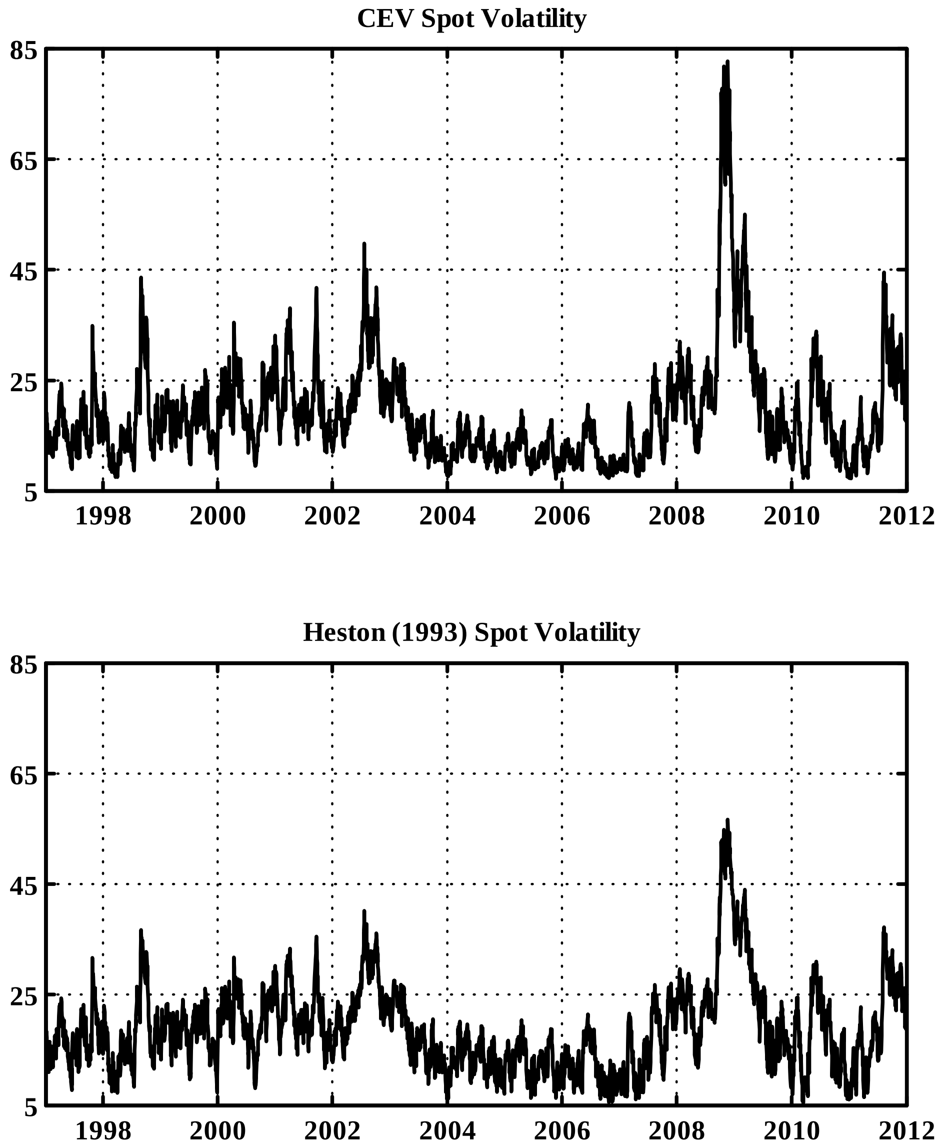

Notes to Figure: The figure plots the daily spot volatilities filtered from S&P 500 daily returns using particle filtering. In the top panel, we plot the spot volatilities estimated for the CEV dynamic. In the bottom panel, we graph the filtered spot volatilities obtained for Heston (1993). For both panels, the daily volatilities are annualized and expressed in percentage.

Figure 5: One-Month Variance Risk Premium, IVRMSE, and Implied Volatility Smile by Model

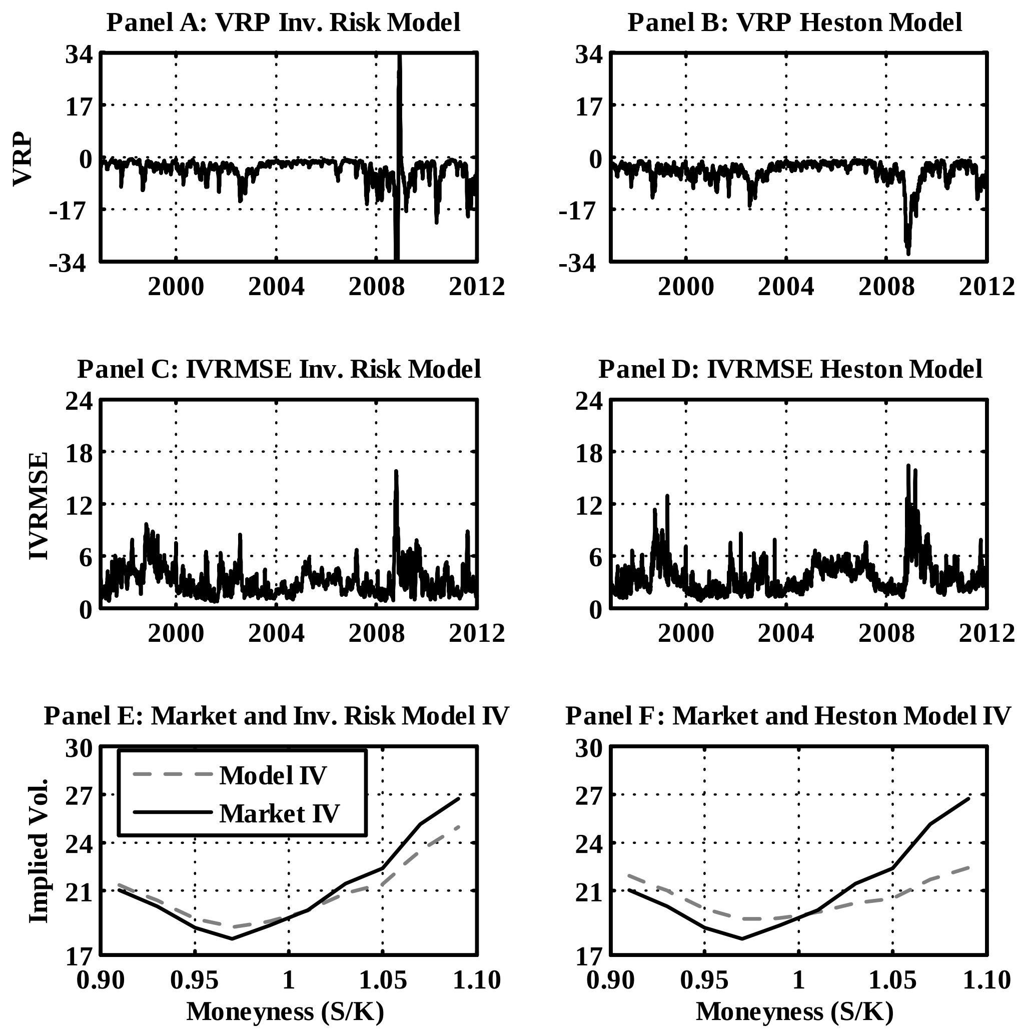

Notes to Figure: For each model, we plot the daily one-month variance risk premium (VRP) and IVRMSE expressed in percentage terms in Panel A to D. To obtain the models' one-month VRP on each day, we simulate 10,000 paths, calculate the 30 days integrated VRP for each path, and take the average. For the IRW model, the VRP is calculated using estimated parameters and latent variables. For the Heston model, the instantaneous VRP is set to $h \cdot V_t$ where $h = -1.08$. In Panel E and F, we graph market-implied (solid) and model-implied (dashed) volatility smile.

Figure 6: The Dollar Response of Index Options to Inventory Risk and Market-Maker's Wealth

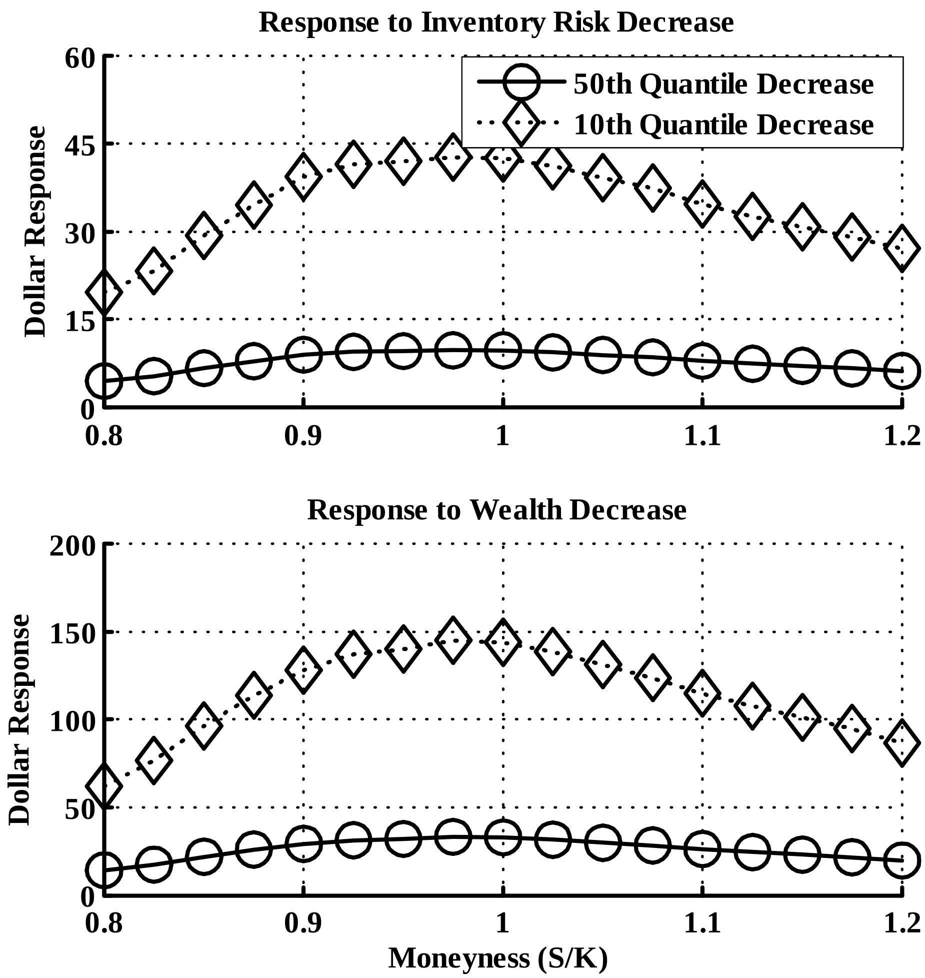

Notes to Figure: We plot the dollar response of SPX puts with 90 days to maturity to a decrease in inventory risk (top panel), and to a decrease in market-maker wealth (bottom panel). To calculate the model sensitivies $\frac{\partial P}{\partial InvRisk}$ and $\frac{\partial P}{\partial W}$, we use the estimated parameters $\hat{\Theta}^{V}$, $\hat{\Theta}^{Inv}$, and $\hat{w}$, and set $r = 4\%$, $q = 0$, $S_t = 1183$, and $V_t = \hat{\theta}$. Moreover, inventory risk and market-maker wealth are set to $InvRisk_t = -9.03E+09$ and $W_t = \hat{w}$. Based on these sensitivities, we then calculate the dollar response of each option as $\Delta InvRisk_t \cdot \frac{\partial P}{\partial InvRisk}$ and $\Delta W_t \cdot \frac{\partial P}{\partial W}$, and plot the result across moneyness. In both panels, each line corresponds to different level of decrease in inventory risk and wealth. The circles correspond to the dollar response to an average decrease in the state variables. The diamonds identify the response when the decrease of each latent variable is within its 90^{th} percentile decrease.

Table 1: Descriptive Statistics

SPX Option Statistics

Year	Implied Volatility (%)	Vega	Days to Maturity	Quotes	Volume	$VRP_{t,30}$ (%)
1997	22.18	133	102	214	303	-1.95
1998	25.14	164	101	181	324	-2.35
1999	24.96	217	110	178	290	-2.60
2000	22.95	223	111	138	273	-0.93
2001	23.93	186	116	159	316	-2.64
2002	25.20	149	109	157	399	-2.45
2003	21.64	140	109	153	503	-2.71
2004	16.44	156	109	186	525	-1.17
2005	14.23	148	95	171	799	-0.63
2006	14.57	151	85	217	1,109	-0.68
2007	18.69	203	101	343	1,073	-1.17
2008	29.39	181	101	457	826	-1.29
2009	29.51	136	98	474	709	-5.61
2010	22.30	160	106	516	708	-2.63
2011	23.04	170	95	545	713	-3.08
Average	22.28	168	103	273	591	-2.13

Note to Table: For each year, we report the average of implied volatility, vega, days to maturity, number of quotes, and volume for SPX puts and calls combined. We also report the average of the one-month variance risk premium measured as the difference between the one-month ex-post realized variance and one-month expected risk-neutral variance. Option implied volatility and vega are computed using Black-Scholes.

Table 2: Implied Volatility, Market-Makers' Inventory, and Delta-Hedged Gains and Losses by Moneyness and Maturity for SPX Options

Days to Maturity		*Moneyness (S/K)* 0.80 to 0.85	0.85 to 0.95	0.95 to 1.05	1.05 to 1.15	1.15 to 1.20	Sum of Inventory by Days-to-Maturity
10 to 30	IV (%)	43.80	21.27	21.00	29.06	38.48	
	Inventory	-1,419	-2,010	-15,100	-1,856	-3,616	*-24,000*
	ΔHedge ($)	14.27	5.68	-8.16	-10.29	-5.87	
31 to 60	IV (%)	27.81	17.71	20.73	26.47	31.36	
	Inventory	-1,226	-4,309	-12,574	-7,897	-777	*-26,784*
	ΔHedge ($)	14.59	2.22	-9.08	-10.92	-11.26	
61 to 90	IV (%)	21.99	17.60	20.83	25.60	29.10	
	Inventory	-499	-1,376	-6,567	-8,672	-946	*-18,061*
	ΔHedge ($)	11.25	2.67	-5.08	-7.39	-13.54	
91 to 120	IV (%)	20.10	18.85	22.33	26.65	29.93	
	Inventory	-332	-1,288	-6,752	-9,241	-2,324	*-19,937*
	ΔHedge ($)	11.18	2.89	-5.24	-7.05	-19.65	
121 to 365	IV (%)	17.52	18.63	21.32	24.22	26.21	
	Inventory	464	1,628	-1,615	-8,900	-2,159	*-10,581*
	ΔHedge ($)	7.14	-0.53	-3.48	-4.65	-14.57	
Sum of Inventory by Moneyness		*-3,012*	*-7,354*	*-42,608*	*-36,566*	*-9,823*	***Total Inventory -99,363***

Note to Table: For each moneyness and days to maturity category, we compute the average for SPX options of implied volatility denoted IV, and market-makers' inventory. We also report ΔHedge, which denotes the sample average of the daily delta-hedged gains and losses across all options in each moneyness and maturity category. The top right column reports the sum of inventory for each days to maturity group while the bottom row reports the sum of inventory by moneyness category. On each day, an estimate of market-makers inventory is obtained by summing inventory across all contracts within the moneyness and maturity category considered.

Table 3: Time Variation in the One-month Log-Variance Risk Premium

	$\Delta LogVRP_{t,30} \times 100$					
	1997-2011		*1997-2004*		*2005-2011*	
	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*
Intercept	-0.26	-0.06	-0.43	-0.23	-0.10	-0.01
	(*0.17*)	(*0.80*)	(*0.11*)	(*0.55*)	(*0.70*)	(*0.98*)
InvRisk(*t-1*)		0.96 ***		0.85 ***		1.20 ***
		(*0.00*)		(*0.00*)		(*0.00*)
ΔW(*t*) × InvRisk(*t-1*)		-3.39 ***		*-5.82* ***		*-5.11* ***
		(*0.00*)		(*0.00*)		(*0.00*)
S&P500LogRet(*t*)	7.53 ***	7.46 ***	7.32 ***	6.84 ***	7.97 ***	7.91 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
JumpFactor(*t*)	*-4.89* ***	*-4.06* ***	-4.12 ***	-2.28 ***	-5.74 ***	-4.06 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
NetBuyingPressure(*t*)	-0.22	-0.14	-0.03	-0.16	-0.62 **	-0.31
	(*0.26*)	(*0.45*)	(*0.92*)	(*0.53*)	(*0.02*)	(*0.17*)
Disagreement(*t*)	*-0.46* ***	*-0.61* ***	0.01	-0.26	-0.44	-0.74 **
	(*0.07*)	(*0.01*)	(*0.98*)	(*0.50*)	(*0.19*)	(*0.03*)
$\Delta LogVRP_{t-1,30}$	-3.25 ***	-2.97 ***	-5.02 ***	-3.94 ***	-0.99 ***	-0.39
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.01*)	(*0.22*)
Adj. R^2 (%)	39	43	34	43	49	57
N	3774	3774	2011	2011	1763	1763

Note to Table: The table presents the full sample and subsample results obtained from regressing the daily changes in the one-month log-variance risk premium denoted $\Delta LogVRP_{t,30}$ on lagged inventory risk, and the interaction of changes in market-makers' wealth with lagged inventory risk. We control for the effects of S&P 500 log-return, S&P500 jump, net buying pressure for index options, and investors' disagreement. On each day, we estimate inventory risk by summing vega-weighted inventories across options. Finally, changes in market-makers' wealth are measured as the sum of delta-hedged inventory profits and losses (2.6) with bid/ask spread revenue (2.7). All regressors are standardized to have unit variance. The *p* -values are in parentheses and are computed using Newey-West with 8 lags.

Table 4: Time Variation in the Term Structure of Log-Variance Risk Premia

	$\Delta LogVRP_{t,T}$ × 100				
	T=30	*T=60*	*T=90*	*T=180*	*T=270*
	(1)	*(2)*	*(3)*	*(4)*	*(5)*
Intercept	-0.06	-0.09	-0.16	-0.03	0.01
	(*0.80*)	(*0.57*)	(*0.21*)	(*0.72*)	(*0.94*)
InvRisk(*t-1*)	0.96 ***	0.64 ***	0.39 ***	0.41 ***	0.25 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
ΔW(*t*) × InvRisk(*t-1*)	-3.39 ***	-2.51 ***	-2.03 ***	-1.46 ***	-1.28 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
S&P500LogRet(*t*)	7.46 ***	6.39 ***	5.58 ***	4.08 ***	3.52 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
JumpFactor(*t*)	-4.06 ***	-2.37 ***	-1.85 ***	-1.10 ***	-0.79 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
NetBuyingPressure(*t*)	-0.14	-0.17	-0.15	-0.07	0.06
	(*0.45*)	(*0.21*)	(*0.17*)	(*0.34*)	(*0.62*)
Disagreement(*t*)	-0.61 ***	-0.45 **	-0.53 ***	-0.33 ***	-0.41 ***
	(*0.01*)	(*0.03*)	(*0.00*)	(*0.00*)	(*0.00*)
$\Delta LogVRP_{t-1,T}$	-2.97 ***	-1.48 ***	-1.20 ***	-0.53 ***	-0.16
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.23*)
Adj. R^2 (%)	43	47	49	53	20
N	3774	3774	3774	3774	3774

Note to Table: The table presents the coefficients obtained from regressing the daily changes in the log-variance risk premium denoted $\Delta LogVRP_{t,T}$ on the explanatory variables. When measuring the log-variance risk premium we consider five horizons T to capture the term structure of variance risk premia. All regressors are standardized to have unit variance. The p-values are in parentheses and are computed using Newey-West with 8 lags. The sample period is 1997-2011.

Table 5: Parameters, *p*-Values, ajusted R-Squared and Variance Risk Premia Statistics from Daily Estimations of HAR-RV Model

Panel A: Parameter Coefficients and Fit for HAR-RV Model

	$\ln(RV_{t,T})$				
	T=30 (1)	*T=60* (2)	*T=90* (3)	*T=180* (4)	*T=270* (5)
Intercept	-4.10 **	-5.12 **	-5.29 **	-4.11 *	-3.75 **
	(*0.03*)	(*0.02*)	(*0.03*)	(*0.01*)	(*0.02*)
$\ln(RV_{t-1,1})$	0.11	0.06	0.04	0.00	0.01
	(*0.13*)	(*0.28*)	(*0.36*)	(*0.48*)	(*0.46*)
$\ln(RV_{t-6,6})$	0.18	0.10	0.06	0.00	0.00
	(*0.26*)	(*0.43*)	(*0.46*)	(*0.48*)	(*0.46*)
$\ln(RV_{t-30,30})$	-0.02	0.08	0.02	0.02	-0.01
	(*0.35*)	(*0.36*)	(*0.33*)	(*0.29*)	(*0.27*)
$\ln(RV_{t-60,60})$	*0.28*	*0.08*	*0.08*	*-0.13*	*-0.04*
	(*0.26*)	(*0.29*)	(*0.29*)	(*0.26*)	(*0.24*)
$\ln(RV_{t-90,90})$	-0.23	-0.07	-0.15	0.08	0.06
	(*0.24*)	(*0.27*)	(*0.28*)	(*0.23*)	(*0.25*)
$\ln(RV_{t-120,120})$	*-0.37*	*-0.58*	*-0.45* *	*-0.10* *	*-0.06* *
	(*0.20*)	(*0.11*)	(*0.10*)	(*0.09*)	(*0.08*)
Adj. R^2 (%)	51	56	60	71	73

Panel B: Min, Max, and Average of Variance Risk Premia Implied by the HAR-RV model

	$VRP_{t,T} \times 100$				
	T=30	*T=60*	*T=90*	*T=180*	*T=270*
Min	-35.48	-30.06	-22.01	-14.72	-18.22
Max	12.16	6.08	16.27	50.02	17.63
Average	-2.38	-2.46	-2.38	-2.14	-2.08

Note to Table: Panel A reports the average of the parameter coefficients, *p*-values and adjusted R-Squared obtained from estimating the HAR-RV model on each day. The *p*-values in parentheses are computed using Newey-West with 8 lags. Using the coefficients estimated daily from rolling-window of 252 observations we forecast future realized variance for each horizon. Based on the model prediction and the expected risk-neutral variance inferred from option prices we construct measures of variance risk premium. We report the summary statistics of the variance risk premia obtained in Panel B. We refer to the Appendix A for additional information about our estimation methodology. The sample period is 1997-2011.

Table 6: Average Parameters and *p*-Values Across Horizons

	$\Delta LogVRP_{t,T} \times 100$		
	$RV_{t,T}$ proxy by Future Realized Variances (1)	*$RV_{t,T}$ proxy by HAR-RV Model Forecast (2)*	*Parameter Difference (1) - (2)*
Intercept	-0.07	-0.10	0.03
	(*0.65*)	(*0.53*)	
InvRisk(*t-1*)	0.53 ***	0.55 ***	-0.02
	(*0.00*)	(*0.00*)	
ΔW(*t*) × InvRisk(*t-1*)	-2.13 ***	-2.28 ***	0.15
	(*0.00*)	(*0.00*)	
S&P500LogRet(*t*)	5.41 ***	5.45 ***	-0.04
	(*0.00*)	(*0.00*)	
JumpFactor(*t*)	*-2.03* ***	*-1.93* ***	-0.10
	(*0.00*)	(*0.00*)	
NetBuyingPressure(*t*)	-0.10	-0.32	0.22
	(*0.36*)	(*0.16*)	
Disagreement(*t*)	*-0.47* ***	*0.18*	-0.65
	(*0.01*)	(*0.35*)	
$\Delta LogVRP_{t-1,T}$	-1.27 **	-2.16 *	0.89
	(*0.05*)	(*0.01*)	
Adj. R^2 (%)	42	44	

Note to Table: The table reports the average parameters and *p*-values from regressing the daily changes in the log-variance risk premium measured for various horizons on the explanatory variables. The first column reports the average parameters and *p*-values from Table 4 across *T*. In the second column, we present the average parameters and *p*-values obtained when the log-variance risk premia are constructed based on the HAR-RV model prediciton for $RV_{t,T}$. To avoid redundancies details on the results obtained for each horizon when using the HAR-RV model for constructing variance risk premia are provided in Table A.1 in the Online Appendix. All regressors are standardized to have unit variance and the *p*-values are computed using Newey-West with 8 lags. The sample period is 1997-2011.

Table 7: Time Variation in the Term Structure of Variance Risk Premia

	$\Delta VRP_{t,T} \times 100$				
	T=30	*T=60*	*T=90*	*T=180*	*T=270*
	(1)	*(2)*	*(3)*	*(4)*	*(5)*
Intercept	0.00	0.00	0.00	0.00	0.02
	(*0.86*)	(*0.81*)	(*0.74*)	(*1.00*)	(*0.38*)
InvRisk(*t-1*)	0.06 ***	0.04 ***	0.02 *	0.02 ***	0.02 ***
	(*0.00*)	(*0.01*)	(*0.07*)	(*0.01*)	(*0.01*)
ΔW(*t*) × InvRisk(*t-1*)	-0.17 ***	-0.11 ***	-0.09 ***	-0.06 ***	-0.06 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
S&P500LogRet(*t*)	0.83 ***	0.66 ***	0.54 ***	0.36 ***	0.30 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
JumpFactor(*t*)	*-0.19* ***	-0.09 ***	-0.06 ***	-0.03 **	-0.01
	(*0.00*)	(*0.00*)	(*0.01*)	(*0.02*)	(*0.32*)
NetBuyingPressure(*t*)	-0.01	-0.02	-0.02 *	-0.01	0.02
	(*0.44*)	(*0.13*)	(*0.08*)	(*0.19*)	(*0.48*)
Disagreement(*t*)	*-0.06* **	-0.05 *	-0.04 **	-0.03 ***	-0.04 **
	(*0.05*)	(*0.06*)	(*0.03*)	(*0.01*)	(*0.05*)
$\Delta VRP_{t-1,T}$	-0.31 ***	-0.15 ***	-0.10 ***	-0.05 ***	0.00
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.37*)
Adj. R^2 (%)	44	50	53	57	4
N	3774	3774	3774	3774	3774

Note to Table: The table presents the coefficients obtained from regressing the daily changes in the variance risk premium denoted $\Delta VRP_{t,T}$ on the explanatory variables. As in Table 4, we consider five horizons T to capture the term structure of variance risk premia. All regressors are standardized to have unit variance. The p-values are in parentheses and are computed using Newey-West with 8 lags.

Table 8: Statistics of the Return Distribution of Delta-Hedged Near-the-Money Options

Period Length	Min (%)	Max (%)	Average Daily Return per Period × 30 (%)	Annualized Volatility (%)	Skewness	Excess Kurtosis
			Full Sample			
Daily	-12.19	35.46	-8.11	71.11	2.33	10.72
Weekly	-3.92	9.02	-8.37	32.65	1.27	3.42
Monthly	-2.08	4.84	-8.21	17.27	1.44	5.21
			August 9th, 2007 - April 2nd, 2009			
Daily	-10.43	35.46	6.26	95.68	2.55	11.47
Weekly	-3.92	7.94	3.31	42.07	1.34	2.93
Monthly	-1.38	4.84	3.48	28.38	1.95	3.58

Note to Table: We report the min, max, average, volatility, skewness, and excess kurtosis of the return distribution of delta-hedged near-the-money options. Every day, we delta-hedged one long position in each option with moneyness between 0.98 and 1.02 assuming daily rebalancing. We then average the daily returns over all the options used on that day to obtain one single return observation. To obtain weekly and monthly return measures, we average the daily returns over each week and month respectively. We report the statistics of the return distribution calculated for the full sample, and the financial crisis separately. For the financial crisis, we focus on the period starting on August 9th, 2007 when BNP Paribas freezes three of their funds due to valuation issues, and ending on April 2nd, 2009 when the G20 agrees on a global stimulus package worth five trillion dollars.

Table 9: Return Statistics and Parameter Coefficients for CEV and Heston (1993) from MLIS

Panel A: Annualized Statistics for Daily S&P 500 Returns

S&P500 Daily Returns	
Mean	Variance
5.860%	4.583%

Panel B: CEV MLIS Results on Daily Returns Using 10,000 Particles

Parameter Estimates, p-values, and Standard Errors					MLIS Objective Value	Filtered Spot Variances	
κ	θ	δ	ρ_V	η		Mean	Variance
2.91	4.56%	98.35%	-0.60	0.90			
(0.002)	(0.000)	(0.012)	(0.000)	(0.000)	11,746.29	4.433%	142.728%
0.922	0.009	0.394	0.129	0.061			

Panel C: Heston (1993) MLIS Results on Daily Returns Using 10,000 Particles

Parameter Estimates, p-values, and Standard Errors					MLIS Objective Value	Filtered Spot Variances	
κ	θ	δ	ρ_V	η		Mean	Variance
5.32	4.08%	18.82%	-0.47	0.50			
(0.000)	(0.000)	(0.014)	(0.002)		11,722.79	4.084%	52.913%
0.701	0.003	0.077	0.152				

Note to Table: Panel A presents the descriptive statistics for the daily S&P 500 returns. Panel B presents the results when the objective variance is defined by equation (3.2). Panel C presents the results for Heston (1993) when the objective variance is defined by equation (3.2) with η = 1/2. For Panels B and C, the structural parameters and daily spot variances are obtained by maximum likelihood importance sampling (MLIS) based on S&P 500 daily returns. For estimation purposes, the difference μ-q is fixed to the sample average 5.86%. All parameters and statistics are in annual units. The p-values are in parentheses below which are their corresponding standard errors computed using the outer product of the gradient evaluated at the optimal parameter values.

Table 10: Inventory Risk and Market-Maker's Wealth Parameters, Heston (1993)'s Variance Risk Premium Coefficient, and SIVSE

Panel A: Inventory Risk and Initial Wealth Parameters, and Model SIVSE Using 10,000 Simulations, and Based on 6,292 Put Options

Inventory Risk and Wealth Parameters						Sum of Implied Volatility Squared Errors
λ	α	ψ	σ	ρ_{Inv}	w ($ Millions)	
10.73	-1.00E+10	2.28E+11	16.55%	-5.14E-04	440.92	6.68

Panel B: Heston (1993) Variance Risk Premium Parameter, and Model SIVSE Using 10,000 Simulations, and Based on 6,292 Put Options

Heston (1993) Parameter	Sum of Implied Volatility Squared Errors
h	
-1.08	7.99

Note to Table: In Panel A, we present the estimates for inventory risk dynamic (3.7) and market-maker initial wealth parameter when the wealth dynamic evolves according to (3.8). To estimate both dynamics, the variance risk premium is equal to the result in Proposition 2, and the state variables evolve according to Proposition 3. Panel B reports the results for Heston (1993) when the variance risk premium is defined as the product of a constant h with the spot variance. For each model, we also report the sum of the implied volatility squared errors based on 6,292 SPX put observations. All parameters are in annual units. For some of the parameters, we use the scientific notation E*+/-n* to denote the power of 10.

Table 11: In-Sample Fit

Panel A: IVRMSE Over Time

Year	IRW Model	Heston Model	Difference
1997-1999	4.317	3.890	0.427
2000-2002	2.613	2.411	0.202
2003-2005	2.545	3.358	-0.813
2006-2008	3.115	4.027	-0.912
2009-2011	3.114	4.120	-1.006
Min	2.545	2.411	-1.006
Max	4.317	4.120	0.427
Average	3.141	3.561	-0.420

Panel B: IVRMSE by Moneyness

Moneyness	IRW Model	Heston Model	Difference
S/K≤0.95	3.666	3.542	0.124
0.95<S/K≤1.05	3.424	3.675	-0.251
S/K>1.05	4.128	4.292	-0.164
Min	3.424	3.542	-0.118
Max	4.128	4.292	-0.164
Average	3.739	3.836	-0.097

Panel C: IVRMSE by Maturity

Month to Maturity	IRW Model	Heston Model	Difference
2 months	3.829	3.912	-0.082
3 months	3.452	3.800	-0.348
6 months	3.417	4.004	-0.586
Min	3.417	3.800	-0.382
Max	3.829	4.004	-0.174
Average	3.566	3.905	-0.339

Note to Table: The table presents the fit obtained by the intermediary and the Heston model based on 131,638 put observations. For each model the structural parametes are set to their values at the optimum. The performance measure is based on the percentage IVRMSE. In Panel A, we present the IVRMSE averaged by year. Panel B reports the performance of each model by moneyness. In Panel C, we present the models' IVRMSE by maturity.

Table 12: Out-of-Sample Fit

Year	*Panel A: IVRMSE Over Time* IRW Model	Heston Model	Difference
1997-1999	4.209	3.736	0.473
2000-2002	2.642	2.468	0.174
2003-2005	2.572	3.390	-0.819
2006-2008	2.928	3.775	-0.847
2009-2011	3.110	3.967	-0.856
Min	2.572	2.468	-0.856
Max	4.209	3.967	0.473
Average	3.092	3.467	-0.375

Moneyness	*Panel B: IVRMSE by Moneyness* IRW Model	Heston Model	Difference
S/K≤0.95	3.275	3.452	-0.176
0.95<S/K≤1.05	3.338	3.614	-0.277
S/K>1.05	3.791	3.979	-0.188
Min	3.275	3.452	-0.176
Max	3.791	3.979	-0.188
Average	3.468	3.682	-0.214

Month to Maturity	*Panel C: IVRMSE by Maturity* IRW Model	Heston Model	Difference
2 months	3.531	3.694	-0.163
3 months	3.301	3.642	-0.341
6 months	3.327	3.898	-0.571
Min	3.301	3.642	-0.341
Max	3.531	3.898	-0.367
Average	3.386	3.745	-0.358

Note to Table: The table presents the out-of-sample fit obtained by the intermediary and the Heston model based on 131,638 put observations. To obtain the out-of-sample performance of each model we set the structural parameters to the values estimated. For the IRW model, we then set the spot variance and inventory risk on any given day to their one-day-ahead forecast given their values on the previous day. Similarly, the spot variance used for the Heston model on any given day is set to the one-day-ahead predicted value taking the spot variance on the day before as given. We compute the model IVRSME based on these predicted values. In Panel A, we present the IVRMSE averaged by year. Panel B reports the performance of each model by moneyness. In Panel C, we present the models' IVRMSE by maturity.

Table A.1: Time Variation in the Term Structure of Log-Variance Risk Premia Implied by the HAR-RV Model

	ΔLogVRP$_{t,T}$ × 100				
	T=30	*T=60*	*T=90*	*T=180*	*T=270*
	(1)	*(2)*	*(3)*	*(4)*	*(5)*
Intercept	-0.14	-0.07	-0.07	-0.12	-0.10
	(*0.60*)	(*0.73*)	(*0.72*)	(*0.34*)	(*0.24*)
InvRisk(*t-1*)	0.92 ***	0.70 ***	0.59 ***	0.31 ***	0.23 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.01*)	(*0.01*)
ΔW(*t*) × InvRisk(*t-1*)	-3.78 ***	-2.71 ***	-2.31 ***	-1.48 ***	-1.13 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
S&P500LogRet(*t*)	7.79 ***	6.39 ***	5.40 ***	4.19 ***	3.47 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
JumpFactor(*t*)	-3.78 ***	-2.24 ***	-1.74 ***	-1.05 ***	-0.84 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
NetBuyingPressure(*t*)	-0.59 ***	-0.50 ***	-0.38 ***	-0.09	-0.04
	(*0.01*)	(*0.00*)	(*0.01*)	(*0.27*)	(*0.49*)
Disagreement(*t*)	0.67 **	0.34 *	0.05	0.04	-0.18 *
	(*0.03*)	(*0.09*)	(*0.83*)	(*0.76*)	(*0.06*)
ΔLogVRP$_{t-1,T}$	-5.06 ***	-2.75 ***	-2.06 ***	-0.65 ***	-0.27 **
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.03*)
Adj. R^2 (%)	40	42	42	46	49
N	3774	3774	3774	3774	3774

Note to Table: The table presents the coefficients obtained from regressing the daily changes in the log-variance risk premium denoted ΔLogVRP$_{t,T}$ on the explanatory variables. To measure expected physical variance, we fit the HAR-RV model on each day as described in Appendix A and use it to forecast future physical variance. Based on the model forecast we construct measures of the log-variance risk premium. We consider five horizons T to capture the term structure of variance risk premia. All regressors are standardized to have unit variance. The *p*-values are in parentheses and are computed using Newey-West with 8 lags. The sample period is 1997-2011.

Table A.2: Time Variation in the Term Structure of Log-Variance Risk Premia When Delta-Hedged Inventory Gains and Losses Are Calculated Using Ask Prices

	$\Delta LogVRP_{t,T} \times 100$				
	T=30	*T=60*	*T=90*	*T=180*	*T=270*
	(1)	*(2)*	*(3)*	*(4)*	*(5)*
Intercept	-0.23	-0.16	-0.13	-0.10	0.03
	(*0.24*)	(*0.21*)	(*0.23*)	(*0.16*)	(*0.84*)
InvRisk(*t-1*)	0.94 ***	0.64 ***	0.39 ***	0.40 ***	0.25 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
ΔW(*t*) × InvRisk(*t-1*)	-3.40 ***	-2.53 ***	-2.06 ***	-1.46 ***	-1.29 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
S&P500LogRet(*t*)	7.43 ***	6.37 ***	5.56 ***	4.06 ***	3.50 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
JumpFactor(*t*)	-4.02 ***	-2.34 ***	-1.83 ***	-1.09 ***	-0.77 ***
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)
NetBuyingPressure(*t*)	-0.15	-0.17	-0.15	-0.07	0.06
	(*0.45*)	(*0.21*)	(*0.17*)	(*0.34*)	(*0.62*)
Disagreement(*t*)	-0.59 ***	-0.44 **	-0.52 ***	-0.32 ***	-0.41 ***
	(*0.02*)	(*0.03*)	(*0.00*)	(*0.00*)	(*0.00*)
$\Delta LogVRP_{t-1,T}$	-2.95 ***	-1.47 ***	-1.19 ***	-0.52 ***	-0.15
	(*0.00*)	(*0.00*)	(*0.00*)	(*0.00*)	(*0.23*)
Adj. R^2 (%)	43	47	49	53	20
N	3774	3774	3774	3774	3774

Note to Table: The table presents the coefficients obtained from regressing the daily changes in the log-variance risk premium denoted $\Delta LogVRP_{t,T}$ on the explanatory variables. In contrast with Table 4, on each day market-makers' delta-hedged inventory profits and losses are calculated using option ask prices. We consider five horizons T to capture the term structure of variance risk premia. All regressors are standardized to have unit variance. The p-values are in parentheses and are computed using Newey-West with 8 lags. The sample period is 1997-2011.