December 1, 2012

What is Risk Neutral Volatility?

by

Stephen Figlewski*

* Professor of Finance
New York University Stern School of Business
44 West 4th Street
New York, NY 10012-1126

email: sfiglews@stern.nyu.edu
tel: 212-998-0712

I thank the Q-Group and the NASDAQ OMX Derivatives Research Project at NYU for research support for this project. Thanks to Rob Capellini for preparing the GARCH forecasts. Valuable comments from Joanne Hill, Larry Harris, and participants at the Q Group 2012 Spring Seminar, the University of Michigan finance seminar, and the ITAM Finance Conference are greatly appreciated.

ABSTRACT

What is Risk Neutral Volatility?

A security's expected payoff under the real world distribution for stock returns includes risk premia to compensate investors for bearing different types of stock market risk. But Black-Scholes and the great majority of derivatives valuation models developed from it produce the same option prices as would be seen under modified probabilities in a world of investors who were indifferent to risk. Implied volatility and other parameters extracted from options market prices embed these modified "risk neutral" probabilities, that combine investors' objective predictions of the real world returns distribution with their risk preferences. Under Black-Scholes assumptions, real world volatility and risk neutral volatility are equal. But Black-Scholes pricing does not hold in the real world because of unhedgeable risks that bear nonzero risk premia, and the risk neutral volatility that goes into option prices is not the market's best estimate of the volatility that will actually occur. This paper explores what factors relating to both forecasting the empirical distribution of future returns and the risk neutralization process go into the market's risk neutral volatility parameter. Daily risk neutral densities are extracted from S&P 500 index options from 1996-2011 using a model-free procedure. We compute both risk neutral volatility and realized volatility from the observation date through option expiration in order to compare the sensitivity of the two volatility measures to a wide range of variables relating to different manifestations of volatility, such as tail risk, and to the risk neutralization process, such as the general level of consumer confidence and the size of recent volatility forecast errors.

Keywords: risk neutral volatility; implied volatility; option pricing; risk aversion

JEL Classification: G13, G12, G14

One of the major insights of the original Black-Scholes (BS) option pricing model was that the right measure of risk in valuing an option was the return volatility of the underlying stock, with no distinction between market risk and idiosyncratic risk. The other factors going into the equation are all observable, but volatility is not. Over time, research has shown that the constant volatility lognormal diffusion returns process assumed by BS is greatly oversimplified. Equity market volatility changes over time; returns are not exactly lognormal, especially in the tails; there appear to be non-diffusive jumps in both returns and volatility; and there are (possibly numerous) unspanned stochastic factors affecting actual options that can be expected to be priced in the market, with risk premia that themselves may vary stochastically over time.

Volatility is not easy to predict accurately, and no matter how an estimate is calculated from past returns data it is common for option prices in the market to differ substantially from the model values computed with that volatility estimate. The solution adopted by market makers and many others has been to extract implied volatilities from option market prices. Implied volatility makes the model consistent with the pricing observed in the market, but it has the serious conceptual shortcoming that implied volatility (IV) requires one to specify "the market's" option pricing formula. Almost invariably the BS model is chosen for this purpose, but the multiplicity of traded options with different strike prices produces a multiplicity of IVs, known as the volatility "smile" or "skew" (for a single maturity) or the volatility "surface" (for the full set of strikes and expiration dates), even though the underlying stock obviously can have only a single volatility.

This approach, known as "practitioner Black-Scholes," ignores the theoretical inconsistency and simply fits a different IV for each option. Often, the next line of research is to try to model the stochastic behavior of the implied volatility surface, despite the highly questionable logic of assuming that dependable insights about real world option pricing can be obtained by modeling the behavior of output from a theoretical model that is visibly inconsistent with the way the market actually prices options.

The BS model has the initially surprising property that, even though the expected return on the underlying stock is the risk free interest rate plus a suitable risk premium and the actual payoffs expected from options written on that stock obviously must depend on the stock's expected return, the option's theoretical value relative to the stock does not depend on the stock's risk premium. The BS option value can be computed as the expected payoff under a lognormal distribution with the same volatility as the stock but with a mean return equal to the riskless rate. This modified probability distribution is called the "risk neutral" distribution.

Harrison and Kreps (1979) proved that in a world free of profitable arbitrage opportunities, there will always exist at least one risk neutral distribution that combines investors' risk preferences with their objective forecast of the probability density for the stock price at option expiration. Security prices in the real world market will be equal to their expected payoffs under the risk neutral distribution discounted back to the present at the riskless interest rate. As valuation models have been developed for more elaborate underlying returns processes, the risk neutral density, which we will refer to as the RND, has become more elaborate as well. It must impound the market's estimate of the empirical or real world probability density (often called the "P measure") and the market's current attitudes about bearing every type of risk that cannot be fully hedged (transforming the P measure into the risk neutral "Q measure").

Implied volatility, and any other parameter that is extracted from option prices in the market, is a risk neutral value under the Q measure, and will rarely equal the real world value under the P measure. Recent research on volatility risk, e.g., Carr and Wu (2008), has concluded that the market places a negative risk premium on volatility risk. The market dislikes volatility, so prices for options and other securities whose payoffs increase with higher volatility will be bid up because of their hedging value. They will be priced (under the Q measure) as if volatility were expected to be higher than it really is. This raises the question that is the object of this paper: "What is Risk Neutral Volatility?"

This is an important issue for understanding how the options market works and how the risk-neutralized factors that determine option prices in the market are related to objective forecasts of their real world values.

There has been an enormous amount of research on implied volatility over the years.[1] More recently, attention has begun to turn toward using option prices to extract not just implied volatility, but the entire risk neutral probability density for the price of the underlying on expiration day.[2] Breeden and Litzenberger (1978) showed how this could be done very simply, given prices for options with a dense set of exercise prices that span a broad range of future stock prices. There are two significant problems in executing the procedure: first, how to smooth and interpolate market option prices to limit the effect of pricing noise and to produce a smooth density function, and second, how to extend the distribution to the tails beyond the range of traded strike prices. But one enormous advantage of the RND is that is does not depend on an assumed pricing model. It is model-free.

This paper applies the Breeden-Litzenberger technology to extract RNDs from daily S&P 500 index option prices and then to compute the standard deviations. I explore whether, and how much, the risk neutral volatility is influenced by a broad set of exogenous factors that I hypothesize may be related to estimation of the empirical density or to the risk neutralization process. The primary objective is not to build a formal model of option risk premia, but to establish a set of stylized facts that any such model will need to be consistent with.

[1] See Poon and Granger (2003) or Jackwerth (2004) for reviews.
[2] See, for example, Bliss and Panigirtzoglou (2002), Giamouridis and Skiadopoulos (2012), Kozhan, et al. (2010), or Christoffersen, et al. (2011).

Several broad questions to be addressed include the following. What return and volatility-related factors are most important to investors in forming the forecast of the empirical probability density that is embedded in the RND? What are the relevant time horizons investors focus on? To what extent is the market forward-looking versus backward-looking in gauging volatility? What factors influence the process of risk neutralization? Do the answers to these questions differ for long maturity vs. short maturity options, or in different market conditions?

Understanding what factors influence the RND and how it behaves are important for a number of reasons. Option parameters implied from market prices are often felt to be more accurate, or at least more meaningful, than those obtained strictly from statistical analysis of historical data. Option pricing entails forecasting the future, not just analyzing historical data, and market participants can have information about important future events that is not reflected in the historical record. But extracting the market's true expectations from market prices requires understanding how the risk-neutralization process works. For example, implied volatility is higher than the market's best estimate of true volatility, as was discussed above and is readily confirmed in the time series: Over the 15-year period that we consider below, the VIX volatility index, one estimate of the risk neutral volatility, averaged 22.1 percent while realized volatility was distinctly lower, averaging 20.6 percent.

The Chicago Board Option Exchange created the VIX index as a measure of implied volatility for the overall stock market portfolio. Originally, the VIX was a weighted combination of Black-Scholes IVs, but in 2003 the calculation procedure was changed (as was the underlying index, with the S&P 500 replacing the S&P 100). The computation of the new VIX is based on roughly the same model-free methodology drawn from Breeden and Litzenberger (1978) that we employ below.[3] But, there are several reasons to prefer analyzing the whole RND rather than the VIX. First, the VIX focuses only on volatility and is designed to produce a single number. This loses a large amount of information about the tails and other aspects of the shape of the market's risk neutralized density function that are captured and reflected in the RND. Second, the VIX calculation ignores several issues of data handling that we try to deal with more carefully, including extending the tails of the density beyond the range of available strike prices, using interpolation to reduce approximation errors when numerical derivatives are computed from options with widely spread strike prices, and smoothing the jump in implied volatilities that occurs in the transition from puts to calls in the VIX formula. Jiang and Tian (2007) demonstrate that these features of the VIX computation methodology introduce substantial and unnecessary inaccuracy. Third, the VIX focuses on a single 30-day maturity, while the options we analyze below have maturities ranging from 14 to 199 days. Finally, while the VIX and the RND volatility we calculate here are not the same, a preliminary investigation showed that replacing RND volatility with VIX in our regressions containing a full set of explanatory variables produced fairly similar results. A full exploration of the differences between the two volatility measures is well worth doing, but it is beyond the scope of this paper.

[3] See Chicago Board Options Exchange (2003).

Black and Scholes assumed the returns distribution came from a constant volatility lognormal diffusion, but this is too simple. Differences that may matter greatly to investors and market makers include tail behavior, time-variation, and other volatility-related properties. Fat tails are important because the empirical distribution has considerably fatter tails than the lognormal allows for, which translates to greater tail risk, but also higher potential returns to options. Asymmetry in the tails matters, too. The lognormal is positively skewed, but the RND is consistently negatively skewed (see Figlewski (2010)), and there is much evidence that investors are especially concerned about downside risk. The response of both risk neutral and empirical volatility following a return shock is asymmetrical between the left and right tails.

Time-variation in the RND is also potentially important to investors, beyond the obviously greater difficulty of hitting a moving target, because stochastic volatility is unspanned (meaning it can't be fully hedged by other securities) and should therefore bear a risk premium (possibly time-varying, as well). Finally, while the single volatility parameter in a lognormal diffusion plays several related roles, this does not hold for other distributions and returns processes. Specifically, in a lognormal diffusion, instantaneous volatility determines the cost and risk of a rebalanced delta hedge, which is what matters to a market maker. The assumption of i.i.d. return shocks means that the 1-day variance is the integrated value of the instantaneous variance, and the total variance over the T-day life of an option is just T times the one-day variance. So the standard deviation of the RND that applies to the possible price of the underlying stock on expiration day is the same parameter as the market maker hedging the option for a very short time wants to know. It is the parameter that governs how common and how extreme returns in the tails of the distribution are. The same volatility determines the probable trading range of the stock over a finite horizon, and the potential profit to an investor who plans to liquidate his option position whenever the stock gets to a predetermined price at any point during its life. With a different underlying returns process (especially once it is modified by risk-neutralization) these various volatility-related concepts may differ from each other, and from the single volatility parameter of a lognormal. Below, we will examine possible explanatory variables designed to proxy for each of these different volatility-related factors.

Finally, since risk premia are significant components of total return, there is considerable value in understanding risk-neutralization for investors and market makers just to be able to estimate the prices of different types of risk. How much is the market paying for risk bearing? Or the flip side, how much does the market charge for hedging those risks? For example, Bollerslev and Todorov (2011) argue that much of the volatility risk premium during the fall of 2008 was compensation for left tail "crash risk," while Kozhan et al. (2010) find that about half of the implied volatility "skew" reflects a skew risk premium.

Before moving on, we must mention the recent theoretical breakthrough in this area by Stephen Ross (2011) and consider how it applies to this research. The state of thought on the issue of recovering both the risk neutral density and the empirical density from market prices for securities has long been that the RND can be extracted from option prices, but objective probability estimates and risk premia cannot be separately deduced from it without further information about either the market's risk preferences or its beliefs about the true returns generating process. Ross proves that this is not strictly true. This is a deep theoretical result with implications for empirical research that have not yet been fully explored. Ross's derivation depends on several critical assumptions, including that markets are complete, which means that there are enough traded securities to span the full set of possible future states of the world, and they can be priced as if all investors were identical to a composite "representative agent." Those strong assumptions cannot be held with confidence as a description of the real world, in which investors have different endowments, different attitudes about risk, different beliefs about the true returns generating process, different exposures to a multiplicity of risk factors they will want to hedge, different institutional constraints on trading, and so on, all of which may evolve over time in ways that are path-dependent. In applying the theoretical Recovery Theorem empirically, it is necessary to add assumptions that restrict the states of the world under consideration, for example to allow focusing only on stock market returns without considering factors from the broader economy. It is also necessary to assume enough temporal stability in the system that data from past periods can be combined in estimating the current state and the state space at future dates. In short, Ross's theoretical result is very important, and may eventually be usefully applied by practitioners even though its underlying assumptions are not satisfied, much as the Black-Scholes model is, but it does not eliminate the immediate problem of separating the effect of risk neutralization from objective volatility estimates in risk neutral densities that we are addressing here.

The next section describes how the RND can be obtained from option prices in theory and sketches the actual steps taken to implement the procedure with real world data. The following section explains the factors that we think could be important enough to investors to influence their demand for options, and defines the variables we use to measure those factors. One issue that arises in forecasting based on historical information is how far back the investors are assumed to look in calculating historical volatility. We first investigate the performance of different lag definitions for a number of variables and also different cutoff probabilities to define whether a given day's return falls into the left or right tail. Section 4 looks at univariate relationships in the form of simple correlations between the volatility measures and the explanatory variables, and then presents the results of a grand regression with all of the variables together. Among the interesting results from this exercise is that several variables that are expected to be correlated with investor risk attitudes do in fact appear to have significant explanatory power for RND volatility.

Section 5 explores the extent to which the RND volatility is forward looking, and over what horizon, because investors successfully incorporate relevant information about the future that is not captured through analyzing historical data. We find that RND volatility contains a statistically significant amount of information about the volatility that will be experienced in the future, particularly over the very short run. It largely dominates historical volatility in explaining future volatility and contributes significant explanatory power in combination with a GARCH model forecast. The best fits in terms of t-statistics and R^2 seem to be for future volatility realized over the next 1-2 weeks, rather than all the way through option expiration. Section 6 explores robustness of the model by fitting it on subsets of the data, according to option maturity and the time period. The final section offers some concluding comments.

Section 2: Computing the Risk Neutral Volatility

In the following, the symbols C, S, X, r, and T all have the standard meanings of option valuation: C = call price; S = time 0 price of the underlying asset; X = exercise price; r = riskless interest rate; T = option expiration date, which is also the time to expiration. We will also use f(x) = the risk neutral probability density function, also denoted RND, and F(x) = the risk neutral distribution function.

The value of a call option is the expected value of its payoff on the expiration date T, discounted back to the present. Under risk neutrality, the expectation is taken with respect to the risk neutral probabilities and discounting is at the risk-free interest rate.

$$C = e^{-rT} \int_X^\infty (S_T - X) f(S_T) dS_T \tag{1}$$

Taking the partial derivative in (1) with respect to the strike price X and solving for the risk neutral distribution F(X) yields:

$$F(X) \quad = \quad e^{rT} \frac{\partial C}{\partial X} + 1 \quad . \tag{2}$$

Taking the derivative with respect to X a second time gives the RND function:

$$f(X) \quad = \quad e^{rT} \frac{\partial^2 C}{\partial X^2} \tag{3}$$

In practice, we approximate the solution to (3) using finite differences. In the market, option prices for a given maturity T are available at discrete exercise prices that can be far apart. To generate smooth densities, we interpolate to obtain option values on a denser set of equally spaced strikes.

Let $\{X_1, X_2, ..., X_N\}$ represent the set of strike prices, ordered from lowest to highest, for which we have simultaneously observed option prices, and denote the price of a call option with strike price X_n as C_n. To estimate the probability in the left tail of the risk neutral distribution up to X_2, we approximate $\frac{\partial C}{\partial X}$ at X_2 and compute

$$F(X_2) \cong e^{rT} \frac{C_3 - C_1}{X_3 - X_1} + 1 \quad (4)$$

The probability in the right tail from X_{N-1} to infinity is approximated by,

$$1 - F(X_{N-1}) \cong 1 - \left(e^{rT} \frac{C_N - C_{N-2}}{X_N - X_{N-2}} + 1 \right) = - e^{rT} \frac{C_N - C_{N-2}}{X_N - X_{N-2}} \quad . \quad (5)$$

The approximate density at a strike X_n, $f(X_n)$, is given by:

$$f(X_n) \cong e^{rT} \frac{C_{n+1} - 2C_n + C_{n-1}}{(\Delta X)^2} \quad (6)$$

The RND aggregates the individual risk neutralized subjective probability beliefs within the investor population. The resulting density is not a simple transformation of the true (but unobservable) distribution of realized returns on the underlying asset as in Black-Scholes, nor does it need to obey any particular probability law. Obtaining a well-behaved RND from market option prices is a nontrivial exercise. There are several key problems that need to be dealt with, and numerous alternative approaches have been explored in the literature.

The variable we wish to explain is risk neutral volatility, computed as the standard deviation of the S&P 500 index on option expiration day under the risk neutral probability density. To be able to combine observations from different time periods with varying numbers of days to expiration into a single regression, this standard deviation is converted into an annualized percentage value, which we will refer to as RND volatility.

The RND is extracted from market option prices using a procedure that produces a smooth curve over the full range of stock prices. The steps in the process are sketched out below. For more details, the interested reader is referred to Figlewski (2010) or Birru and Figlewski (2012).

The first step is to obtain synchronous option prices over the widest available range of exercise prices. Since they are being used to calculate the RND at a single point in time, the prices must be as close to simultaneous as possible. Transactions for away from the money contracts even in this highly active market are not frequent enough for this purpose, so it is important to use market maker quotes, which are updated continuously for all traded contracts. We use the daily closing bid and ask prices from the OptionMetrics database, available from WRDS (Wharton Research Data Services). Various screens are applied to eliminate data points whose informational value is questionable, as detailed below.

The prices of the options used can differ by up to two orders of magnitude between deep out of the money and deep in the money contracts. To facilitate interpolating to construct a dense set of option prices, the Black-Scholes equation is used as a kind of transform, not unlike taking logs, to convert from strike price-option price space into strike price-implied volatility space. After the interpolation step, the inverse transformation is applied to return to the original strike price-option price space to calculate the numerical derivatives.

It has become common practice to compute the RND by combining quotes from out of the money puts and out of the money calls because a major component of an in the money option's price is its intrinsic value, which does not provide information about probabilities.[4] This involves joining the two data sets at some strike level, but because BS implied volatilities for puts and calls with the same exercise price are consistently different from one another, to avoid an artificial distortion in the RND at this point, we introduce a smooth transition from puts to calls within a narrow price range around the at the money forward index level.[5,6]

Interpolation is done by fitting a fourth order spline curve to the midpoints of the implied volatilities computed from the option bid and ask price quotes. Earlier papers have used cubic splines in the interpolation step, but the resulting curves are actually not smooth enough. Once the second partial derivative is taken the density fitted from a cubic spline is continuous, but it typically has spikes because its first derivative is not smooth. This problem disappears with a fourth degree spline. The bid-ask spread is quite wide, even in this active options market. We use the Matlab function ***spap2*** to fit fourth order splines with one knot point and then apply the function ***newknt*** to optimize the placement of the knot. In some cases, the fairly large tick size in the options market causes sharp discontinuities in the volatility smile that interfere with fitting a sensible spline approximation. When the "optionality" component of an option's value is small, there can be a very large difference in implied volatility between rounding the price up to the next higher tick versus rounding down. This problem is most severe for deep in or out of the money options with short maturity.[7]

The interpolated IVs are then converted back into option prices and numerical partial derivatives are calculated using eqs. (4)-(6). We now have the portion of the RND that lies between the lowest and highest exercise prices in the sample, that we will refer to as the empirical portion of the RND, as well as the total probability in each of the missing tails.

The remaining portions of the RND are completed by appending tails from a Generalized Pareto Distribution, fitted so that each tail contains the correct total probability mass and the new tails have the same density as the available portion of the empirical RND at two points: where the GPD tail begins and 3 percentage points further into the tail. Where the data permit, the tail connection points are set at the 5th and 95th percentiles, with density matching at the 5th, 2nd,

[4] See Chicago Board Options Exchange (2003) or Bliss and Panigirtzoglou (2002).
[5] This difference between put and call IVs implies a violation of put-call parity. However, transactions costs and risk to execute the arbitrage trade are significant, which allows a sizable disparity to persist.
[6] Specifically, in a range of $F_t \exp(-.05*T^{.5}) < X < F_t \exp(.05*T^{.5})$, where F_t is the forward index level and T is the time to option maturity in years, we use a smooth linear transition from put IV to call IV in the interpolation step.
[7] Note that this problem does not arise because of the conversion from prices to IVs. It is inherent in the fact that rounding the price to the nearest tick in the market changes the option's time value by a large amount in probability terms. Options for which this problem led to clearly erroneous densities, e.g., tail probabilities that were negative, were removed from the sample.

95th and 98th percentiles. If the empirical RND does not extent as far as the 5th percentile, the left tail density is matched at the lowest available strike and the GPD tail is appended at a point 3% higher. Equivalent treatment is applied when the empirical RND's right tail is short. In a substantial number of cases, large portions of one or both tails were missing from the available option quote data. A date/maturity pair was discarded if the empirical portion of the RND did not extend at least to the 3rd percentile on the left or the 93th percentile on the right.[8]

Figure 1 displays three RNDs that were constructed from December maturity option quotes in early October of 2006, 2007 and 2008. They have the same number of days to expiration, but they are very different from one another, reflecting very different economic environments for these years. October 2006 was one of the lowest volatility periods for the stock market in recent memory and the RND for that period is extremely narrow. By the fall of 2007, the volatility had reverted to a more normal level, but this corresponded to the end of a long bull market, with the S&P index at about its highest level in history. The long left tail may be an indication that investors were becoming worried about a sharp drop in the market. October 2008 was right at one of the most volatile periods in history. The S&P had fallen 37% from the previous year and it would drop another 7.6% the next day. Risk neutral volatility was about 45% and rising sharply. It topped out at over 100% later in the fall.

Section 3: Data

This section discusses a broad set of factors that could influence the volatility forecasts or risk attitudes of investors and/or market makers and the variables we have constructed to measure them. The variables are divided into groups: volatility-related dependent variables, explanatory variables relating to information that is available on the observation date, explanatory variables relating to information that must be constructed from historical returns, and explanatory variables relating to risk attitudes.

Data for computing RND and Realized volatility:

The data required for computing the RND as described above were downloaded from OptionMetrics, including closing bid and ask prices for all traded S&P 500 index options that satisfied the following criteria:

[8] The asymmetry in these cutoff values is due to the availability of traded strikes. In a rising market, more strikes will have been introduced in the past at levels that are now further in the left tail, than the newly introduced strikes at higher levels are in the right tail. The opposite problem occurs after the market has fallen sharply, as happened in the fall of 2008.

- The observation date and option expiration were both between Jan. 4, 1996 and April 29, 2011.
- OptionMetrics was able to compute an IV > 0 for the option (meaning the midquote price did not violate static arbitrage bounds).
- The bid price was no less than 0.50. This condition eliminates very deep out of the money options that are highly illiquid (the average index level in the sample is 1151).
- Days to expiration were from 14 to 199 calendar days. This allowed extraction of about three RNDs with different maturities per day.
- Implied volatility curves were sufficiently complete and regular that a 4th order spline could be fitted.
- Sufficient options data were available for the empirical RND from the market to cover at least the 3rd through the 93rd percentiles.

<u>Other data needed for computing RND volatility</u>: Riskless interest rates for a wide range of maturities were obtained from OptionMetrics. The riskless rate was interpolated to match the option maturity. Dividend yields for the S&P 500 index were also downloaded from OptionMetrics. Option maturity dates for most contracts were adjusted: Regular contracts technically expire on the Saturday following the third Friday of the expiration month, but the payoff is fixed at Friday's market opening, so the last time the contract can actually be traded is at Thursday's close. The effective expirations for those contracts were changed to the previous Thursdays. The recently introduced End of Month contracts expire at the close on the last business day of the relevant month, however, so no adjustment to their maturities was needed.

<u>Realized volatility to expiration</u>: Theoretical option prices are based on the cost of replicating the option's payoff by a continuously rebalanced portfolio of the underlying index plus borrowing or lending of the riskless asset. This cost is a function of volatility, so the density impounded in option prices should be the market's best forecast of the volatility that will be realized from the present through option expiration. In the BS world, there is no risk premium on volatility, so volatility under the RND and under the empirical density are the same. We want to explore how the two differ in our world, by comparing regressions of RND volatility on explanatory variables against the same regressions run on realized volatility.

Realized volatility on date t for an option maturing on date T is defined as

$$\sigma_{Realized} = \sqrt{\frac{252}{(T-t)} \sum_{t+1}^{T} r_\tau^2} \qquad (7)$$

This assumes 252 trading days in a year, and adopts the standard approach of treating the daily mean return as equal to 0.[9] Note that Realized volatility here is expressed as an annualized percent, e.g., 20% is 20.0.

We consider explanatory variables in three sets. First are variables related to volatility and extreme price movement, that do not involve analyzing historical data over an unknown sample horizon. In the second set are volatility-related variables that are calculated from past historical data, for which we must decide what lag to use. The third set of variables are chosen to reflect factors that might affect the risk neutralization process.

Date t Volatility Variables

Date t return: It is well-established that volatility, both empirical and implied, has a strong tendency to go up when the stock index falls. Often called the "leverage effect," this phenomenon is offered as one of the primary reasons for the existence of the implied volatility smile pattern in the options market. The GARCH model that we use here includes an asymmetry term to capture the greater effect of a negative return than a positive return on subsequent volatility. Stochastic volatility models, that introduce separate stochastic factors for returns and volatility, typically estimate the two shocks to be negatively correlated, with $\rho \cong -0.7$. The Date t return is defined as 100 x $\log(S_t / S_{t-1})$.

Absolute value of date t return: A large return of either sign is consistent with high volatility. An increase in expected future volatility following a large shock is also consistent with the assumed dynamics of volatility in GARCH and stochastic volatility models. This variable was originally included in the analysis, but it was subsequently replaced with a different but related one, as described below.

GARCH model forecast: Probably the most widely used model for time-varying volatility in the stock market is Generalized Autoregressive Conditional Heteroskedasticity (GARCH). The specific form of the model we use is the Glosten-Jagannathan-Runkle (1993) version (GJR-GARCH), which can be written as

$$r_t \;=\; \sigma_t \, z_t \,, \qquad z_t \sim N(0,1)$$

$$\sigma_t^2 \;=\; a \,+\, b\,\sigma_{t-1}^2 \,+\, c\,r_{t-1}^2 \,+\, d\,1_{\{r_{t-1}<0\}}\,r_{t-1}^2 \qquad (8)$$

where r_t is the date t log return, σ_t^2 is the conditional variance at date t, and z_t is a standard normal i.i.d. random variable. The return equation contains no constant term, consistent with the

[9] The average daily return is only a few basis points, which is much smaller than the average sampling error on the mean, so it is generally felt that the small bias introduced by treating the mean as zero is preferable to the errors that would result from using the sample mean in the calculation.

way realized volatility is defined in equation (7). The variance equation models σ_t^2 as a linear function of the date t-1 conditional variance, yesterday's squared return shock, and an asymmetry term that increases σ_t^2 if the most recent return was negative. The model as shown projects variance one day ahead. Multiday forecasts, as needed to forecast volatility from date t through option expiration, require substituting forecasted values based only on data known at date t for future variances. A key issue is that in fitting forecasting models using historical data, the GARCH model coefficients must be those that would have been available at the time, which have to be estimated using only historical data up to that date.

We use volatility forecasts obtained from Robert Engle's VLAB.[10] A GJR-GARCH model was fitted on past returns data for every day during the sample, and out of sample variance forecasts were computed for 1 to 252 trading days ahead. These were used to construct predictions of average variance from each observation date to option expiration. For easier comparison with other variables, the resulting variances were converted to annualized percent volatilities.

Date t trading range: The intraday trading range is important to option traders for two reasons. First, it contains a substantial amount of information about volatility. For example, Parkinson (1980) showed that a volatility estimator based on the intraday high and low can achieve the same accuracy as one using only close to close returns with five times as many observations. Second, option market makers manage their risk exposure by hedging the "Greek letter risks" of their inventory positions. The hedges need to be rebalanced as the price of the underlying moves, so hedging costs can be expected to be positively related to the size of the range over which the stock price travels in a day. The Date t range is defined as the day's high minus the low as a percent of the average. For consistency with the other variable definitions, this is converted into a logarithmic percentage in the following way:

$$\text{Range}(t) = 100 \times \log(1 + (S_{t,HI} - S_{t,LOW}) / ((S_{t,HI} + S_{t,LOW})/2))\qquad (9)$$

Like the absolute return for the day, this variable was included in our initial exploration, but discussions with practitioners suggested a better way to measure the difficulty in managing a market maker's risk on a day with a large price move is the trading range minus the absolute price change. It is easier to manage risk on a day with a wide trading range if the market simply moves strongly from one level to another, than it is on a day of high uncertainty, when the market fluctuates over a broad range without finding a new equilibrium much different from the initial level.

Date t trading range minus absolute return: In the previous version of this paper, we included both that absolute return and the trading range, as described above. The trading range came in with a highly significant positive coefficient but the coefficient on absolute return consistently was strongly negative and about the same size. Taking the difference between the two, as

[10] http://vlab.stern.nyu.edu/

suggested in discussion with practitioners, produced a combined variable with a strongly positive and significant coefficient.

Date t left and right tail returns: One of the most apparent differences between the empirical distribution of returns and the lognormal is in the tails. Large absolute returns, especially on the downside, are substantially more common in practice than the lognormal allows for, so it is possible that investors and market makers may react to a large tail event differently than to a normal sized return shock. But what constitutes a tail event? Is it a return more extreme than is expected 5% of the time, or 1%? I treat this as an empirical question that I explore by considering the explanatory power of several alternative tail cutoff definitions, both in univariate regressions of RND volatility on a tail variable, and also along with our other explanatory variables in multivariate specifications. From the results of this investigation, which is described in detail below, I define the return on date t as a left (right) tail event if it falls in the 2% left (right) tail as defined relative to the normal. The left (right) cutoff is set at -2.054 (+2.054) times the previous day's 1-day GARCH volatility forecast. If the return is not in the tail, the value of the tail variable is set to 0.

Historical Volatility Variables

Many of the obvious factors that investors and market makers might want to take into account in forecasting future volatility, such as past realized volatility, are computed from the historical return series over some horizon. It is not known a priori what horizon the market focuses on, but the behavior of the RND provides a window into this issue. I explored it by fitting regressions of RND volatility on each variable individually over historical samples of different length and examining the t-statistics and R^2. This showed which horizons gave the closest univariate fits. But, as will be seen below, the coefficients can change sharply in magnitude, sign and significance when combined with other variables in a more realistic specification, so I also looked at regressions with all of the exogenous variables, which gave a different perspective.

Most of the variables involved showed quite similar performance over a fairly broad range of sample lengths, but also some substantial differences with the different metrics. In the end, I selected estimation windows that seemed reasonable in terms of performance in the various exploratory regressions and that "made sense" in terms of the behavior one might expect from an intelligent investor.

Defining left and right tail events: We first consider how to set the cutoff values in defining tail events. In addition to the sizes of left and right tails on date t, we also want to include data on tail events in the recent past. The empirical distribution's fat tails can be manifested as more frequent outcomes in the tails than would be predicted from a normal distribution and/or unexpectedly large tail outcomes conditional on falling in the tail. This distinction is equivalent to the difference between the standard risk measures of Value-at-Risk and Expected Shortfall. For example, the 5th percentile (e.g., the 5% Value at Risk) may be quite similar for a fat-tailed density and a normal density with the same volatility, while the mean return conditional on being in the tail (the Expected Shortfall) for the two can be much different. We consider both the number and average size of left and right tail events observed in the recent past. This requires specifying both the cutoff values that define the tails and the length of the historical period investors are assumed to consider.

Table 1 reports t-statistics on four tail variables from regressing risk neutral volatility on each one individually, in Panel A, and from multivariate regressions with our other variables, in Panel B using cutoffs of 0.5% to 5.0% and historical samples from 2 days to 2 years (500 trading days), in addition to the current date t. Also, as an alternative to a fixed horizon for RNDs extracted from options of all maturities, we considered setting the size of the historical sample in each case equal to the option's remaining lifetime, for example, explaining RND volatility drawn from 4 month options by considering tail events over the previous 4 months. The t-statistics reported in Table 1 have been corrected for the cross-correlation in the residuals due to overlapping forecast horizons, as will be described in detail below.

From theory and empirical observation, we expect that frequent left tail events will increase risk neutral volatility, so the regression coefficient on the number of tail events and its t-statistic should be positive. A left tail return is negative, so a negative coefficient and t-statistic on its average size correspond to an increase in RND volatility. This reasoning becomes ambiguous for right tail events, however. Frequent large positive shocks to returns may indicate that volatility is high. But if this means the market is rising, it may result in lower future realized and implied volatility. Positive coefficients and t-statistics are consistent with the former interpretation.

Consider first the number of left tail events in the recent past. If the tail cutoff is set at 0.5%, the univariate results in Panel A of Table 1 show negative and insignificant t-statistics for most horizons. In Panel B, these negative coefficients become highly significant. But negative coefficients imply, counterintuitively, that frequent left tail events reduce the volatility investors embed in option prices. Comparing across tail cutoffs and horizons, we see that for each tail cutoff, the t-statistics tend to rise and become significantly positive for longer horizons. And for a given horizon, coefficients mostly rise and become significantly positive for higher cutoff values.

Making the same comparisons between Panels A and B across tail cutoffs and horizons for the average left tail return, we see highly significant negative coefficients in most of the univariate regressions, but significantly positive coefficients in the multivariate specification except for the longest horizons.

Any choice we make is inherently somewhat subjective and arbitrary. Given the results in Table 1, and preferring to use the same historical sample for counting left tail returns and measuring their size, I selected the 2.0% critical value to define the left tail and a historical sample of the past 500 trading days.

For right tail events, most of the coefficients and t-statistics are positive and many are highly significant, especially in Panel A. Although other choices could be justified, I selected the same 2.0% cutoff and 500-day sample period for the right tail as I did for the left tail.

To summarize: Each day in a 500-day historical window ending at date t-1 is treated as described above for the Date t left and right tail return variables, by comparing that day's observed return against tail cutoffs defined as the 2nd and 98th percentiles of a standard normal distribution with variance equal to the previous day's 1-day ahead GARCH forecast. The number of tail observations and their average size is computed for both left and right tails.

We now consider the choice of sample length for the other variables that are computed from past returns. Table 2 presents t-statistics in Panel A and R-squared values in Panel B, from univariate and multivariate regressions for each variable.

<u>Return and Absolute return from date t-n to t-1</u>: The return on the index over the recent past is included for the same reason as the Date t return is included: future realized and implied volatility both go up when the market is down. The return variable is defined as 100 x log(S_{t-1} / S_{t-n}). In Table 2, the t-statistics for this variable are negative, as expected, and highly significant at nearly all sample lengths, with the greatest significance in the multivariate specification at lags of only a few days. The absolute value of the return behaved in a fairly similar way, although lags beyond 25 days showed a drop off in significance. Comparing the R-squared results in Panel B with those in Panel A, the correct lag choice for these variables is less obvious. I was concerned, however, that using returns over a long period like 2 years might simply pick up explanatory power from the fact that the stock market over this period experienced long "bull" and "bear" market conditions with associated differences in volatility. So I selected a short 5-day window for these two return variables.

<u>Historical volatility</u>: In univariate regressions, historical volatility is highly significant no matter how long or short the sample period, but in the multivariate specification, very short horizons yield negative and insignificant coefficients. A 65-day window appears to give close to the best fit for both types of regressions. This is not surprising: 3-month volatility is a common choice among practitioners. For this reason, I chose a fixed 65-day horizon rather than using a window equal to the number of days to option maturity, even though the t-statistics were higher for the latter. This choice is supported by the fact that the regression R-squared statistics are very close for the 65-day and option lifetime regressions even though the t-statistics exhibit greater differences. Historical volatility is computed from returns over the last 65 trading days, treating

the expected drift as zero, as is typically done in computing volatility with daily stock returns. The result is converted to an equivalent annualized volatility. Specifically,

$$\sigma_{Hist} = \sqrt{\frac{252}{65}\sum_{t-1}^{t-65} r_\tau^2} \qquad (10)$$

Average daily range t-n to t-1: The daily range is a proxy for potential hedging costs. Each day's range is computed in the same way as described above for the Date t trading range. The results in Table 2 suggest that averaging the daily range over a historical sample equal in length to the option's remaining lifetime performs distinctly better than any fixed horizon.

Range over t-n to t-1: An investor contemplating unwinding his position early would like to know how far its market price might move between now and expiration, which could allow him to make a good profit even on an option that ultimately expires out of the money. How wide a trading range the underlying index has traversed in the recent past might be a reasonable way to gauge this possibility. The variable is defined as

$$\text{Past trading range} = 100 \times \log(1 + (S_{HI} - S_{LOW}) / ((S_{HI} + S_{LOW}) /2))) \qquad (11)$$

where S_{HI} and S_{LOW} are the high and low values for SPX over a period ending on date t-1. In Table 2 the univariate t-statistic is greatest at 65 days, while the largest t-statistic in the multivariate specification is at 10 days. I selected the intermediate value of 25 days for this variable.

The remaining two variables derived from historical returns data relate to the accuracy of the GARCH model we use. I consider them as measures of potential model risk, as is discussed in more detail below. A positive coefficient on the GARCH error means that RND volatility is higher when the GARCH model has been underpredicting realized variance, so the negative and often significant values in Panel A are problematical. Also, I would like to use the same horizon for both GARCH variables, so I chose an historical sample length equal to the option's remaining lifetime, which yields high t-statistics in the univariate regressions, a positive though insignificant value in the multivariate GARCH RMSE regression, and one anomalous negative coefficient.

Risk Attitude Variables

The variables discussed so far all are tied to the volatility and related characteristics of the empirical returns distribution. The remaining factors we consider are related to the risk neutralization process. That is, they reflect things that investors and market makers might consider in transforming objective measures of risk exposure into market prices for options that are exposed to that risk. There are many such factors that could possibly influence investors. I have selected a few that show promise, but there are many others that have not been examined.

University of Michigan Index of Consumer Sentiment: This is a classic survey of consumers' confidence level. Although it does not apply specifically to the financial markets, it is widely reported in the news media and followed by investors. The monthly series is converted to daily values by simply taking the same figure for every day within the month.

Baker-Wurgler Index of Investor Sentiment: Baker and Wurgler (2006) devised a measure of confidence for investors by combining a number of factors such as the discount on closed-end mutual funds, the first day returns on IPOs, etc., and taking the first principal component. Like the other monthly series, we simply repeat the value for every day within the month.

The Price/Earnings Ratio for the S&P 500 index portfolio: This is a more direct measure of how attractive investors are finding the stock market at a given point in time. Optimistic investors pay higher prices for stocks with a given level of current earnings when they anticipate favorable conditions and strong earnings growth in the future. We hypothesize that, other things equal, optimistic investors will also accept lower risk premia for bearing volatility risk. The monthly P/E ratio for the S&P index was downloaded from Robert Shiller's "Irrational Exuberance" website and converted into daily figures by duplicating the monthly number for each day.[11]

The Baa-Aaa Corporate Bond Yield Spread: This daily credit spread series is the difference between the average yield to maturity on corporate bonds rated Baa by Moody's (the lowest investment grade rating) and bonds rated Aaa (the highest grade). It is a measure of the default risk premium for high grade corporate debt. The series was downloaded from the St. Louis Federal Reserve Bank's FRED database.[12]

GARCH model error over a recent period equal to the option's remaining lifespan: Estimates of theoretical option values and Greek letter risks require input of one or more volatility parameters. It is reasonable to think that investor and market maker confidence in their forecasts of future volatility would depend on how accurate those forecasts have been recently. We compute GARCH model errors by looking backward from date t over a period equal in length to the option's remaining lifetime, computing each day's squared return minus the GARCH model variance for that day and averaging. This variable may be positive or negative, corresponding to average under- and over-prediction, respectively, of the squared returns.

GARCH Model RMSE: This is the root mean squared error of the GARCH forecast described above. This variable measures the impact of prediction errors regardless of their sign, while the previous variable allows for possible asymmetric response to under- versus over-estimating volatility risk.

RND Volatility Premium on Date t-1: To the extent that "market sentiment" may depend on exogenous factors that have not been included, or may simply be persistent for whatever reason,

[11] http://www.econ.yale.edu/~shiller/data.htm
[12] http://research.stlouisfed.org/fred2/

that dependence can be captured by including the previous day's volatility risk premium, which we measure as the difference between the date t-1 RND volatility and the GARCH model forecast of volatility through option expiration for that day. Since the objective is to explain the determinants of risk neutral volatility, simply "explaining" it by its own lagged value is not very revealing, so this variable is not included in the standard "all variables" specification we use below. It is useful, however, to see whether yesterday's volatility risk premium contains information for predicting realized volatility, and also to see how the coefficients on the other variables are affected if this variable is included in the regression.

Section 4. Results

This section reports results from univariate correlations among the volatility-related dependent variables and each of the explanatory variables. We then run grand regressions combining all of the exogenous variables into a single specification to gauge the marginal explanatory power of each factor when included with all of the others.

Correlations

Table 3 shows the grand correlation matrix for the two dependent and all exogenous variables. We consider more than 20 explanatory variables, many of which are meant to cover different aspects of risk, as described above. High correlations are to be expected in some cases, with the attendant problem of possible multicollinearity and difficulty is estimating coefficients cleanly.

Consider first the correlations among the exogenous variables. There are actually only a few correlations above 0.7, and these tend to be where one would expect them. For example, quite a few of the variables are highly correlated with the GARCH volatility forecast, including the Date t trading range, Historical volatility, the Average daily range over the recent past, and the GARCH model RMSE. This is not at all surprising. The GARCH volatility forecast is one of the most important explanatory variables. However, the correlations shown in Table 3 are not so high that precise coefficient estimates can't be obtained, as will be clear below.

The first two columns in Table 3 show the simple correlations between the two dependent variables and the explanatory variables. These correlations indicate the sign and degree of explanatory power that each of these variables would have in a one-variable regression with that dependent variable. A positive correlation implies the regression coefficient would be positive and the R^2 would equal the correlation squared.

The first thing to notice is that the correlation between RND volatility and Realized volatility to expiration is 0.613, well under 1.0. This reflects two things. One is that future volatility is hard to forecast accurately (note that the GARCH forecast correlation with realized volatility is only 0.678 and historical volatility's correlation is 0.589. In addition, RND volatility includes the effects of risk neutralization, which evidently entails substantial modification of the market's best prediction of the empirical density.

Proceeding down columns 1 and 2, the first section contains the Date t volatility variables. The Date t return is slightly negatively correlated with the two dependent variables, which is consistent with the common finding that a negative return increases volatility, but the effect seems quite weak. This may be partly due to the fact that the table shows correlations among levels, not changes. A large negative return might well cause volatility to increase even if there is little correlation between returns and the level of volatility. The absolute value of Date t return is distinctly more important than the signed return, especially for RND volatility, which shows correlation of 0.488.

The GARCH model forecast of volatility from date t+1 to expiration is an extremely important variable. Its correlation with both RND and Realized volatility is close to the largest for any of the exogenous variables. Although the model must be fitted with historical data, GARCH has been placed among the Date t volatility variables because ambiguity over how far back the historical data sample should go is not a major issue with GARCH.

Extreme negative and positive "tail" returns show relatively small correlations, but with the expected signs. A large negative left tail return is associated with higher volatility and so is a positive return in the 2% right tail.

The next section of the table contains the volatility-related variables that are computed from the historical price/return series. With negative correlations on the last week's return and positive correlations on the absolute return, if the market went up (down) over the past week, volatilities narrowed (increased), but volatility increased after a week with a large return of either sign.

Historical volatility is likely to be very similar to GARCH but less accurate as a forecast of future volatility, since it uses the same past returns data, but in a less sophisticated way. Still, it may well be more visible and accessible to many market participants than a full-fledged GARCH model. As expected, the correlations between Historical volatility over the last 65 days and the two dependent variables are very high, but a little lower than for Realized volatility and the GARCH forecast.

Looking backward, two different trading range concepts could be important. If the index moves over a broad range intraday it signals high volatility and also greater hedging risk for market makers who are exposed to trading costs and tracking error in maintaining a delta-hedged position. The average daily range over a number of days equal to the option's remaining lifetime measures that effect, and the correlations with the two volatility variables are strongly positive, more so for risk neutral volatility.

The second trading range concept is of more interest to a longer term investor. A wider range is associated with higher volatility, but on the other hand, if the underlying trades over a wide range during the option's lifetime, it can present the investor with an opportunity to realize a profit in the short run even if the option in question eventually expires out of the money, which may make it more attractive. While RND volatility is very highly correlated with both range variables, they have less explanatory power for the volatility that is actually realized.

The next four variables relate to the prevalence and size of past extreme events. We expect the occurrence of a large number of tail events to increase investors' perceptions of volatility risk and perhaps the degree of their aversion to bearing tail risk, as well, so correlations of volatilities to the number of both left and right 2% tail events should be positive, as the table shows. How big the return is when it falls in the tail is another dimension of tail risk, with unusually extreme returns expected to increase future volatility beyond the effect of the simple occurrence of a tail event. The average left tail return is negative, so the negative correlation in the table means volatility increases when a big negative left tail return happens. The right tail return shows the expected positive correlation. The Realized volatility is also related to these tail variables but with substantially lower correlations in each case.

The last section of Table 3 contains the variables that we expect to be more closely related to the risk neutralization process than to forecasting the empirical density. The University of Michigan Survey of Consumer Sentiment and the P/E ratio for the S&P index are negatively correlated with both RND and Realized volatility. The Baker-Wurgler measure, however, is less correlated with RND volatility and it has the "wrong" sign for Realized volatility. To the extent that these variables can be considered proxies for investor confidence, a negative correlation with RND volatility is expected. The Baa - Aaa yield spread in the corporate bond market is a measure of market-perceived default risk. The variable is strongly positively correlated with the two volatility variables, which suggests that wider credit spreads in the bond market correspond to higher volatility in equity options.

The next two variables measure how accurate the GARCH model forecast has been in the recent past. As described above, the data sample for these two variables covers the most recent period of the same length as the option's remaining life. The GARCH model RMSE is strongly correlated with the width of the RND and with future Realized volatility. Investors perceive more volatility risk, and there seems to be more actual volatility, following a period in which volatility models exhibited large errors, especially if the GARCH model had underpredicted the volatility during the sample period (i.e., mean GARCH error was positive).

The final variable in the table is the previous day's RND volatility premium. There is almost no correlation between yesterday's volatility risk premium and today's RND volatility level (although the premium itself is highly autocorrelated), and the correlation with Realized volatility is actually negative. Thus the difference between risk neutral volatility and a good forecast of empirical volatility does not appear to contain information about the volatility that will be realized in the future or even about tomorrow's risk neutral volatility.

Correcting for Cross-Correlation in the Residuals

The simple correlations in Table 3 show that the factors that were expected to influence RND volatility and Realized volatility appear to do so with the anticipated sign. But correlation among the explanatory variables is high, so it is not clear how much any individual factor will contribute beyond the fact that all of them are correlated with "volatility" broadly defined. To gauge the marginal influence of each variable, Table 4 combines them all into a single grand regression.

Cross correlation among the residuals in this regression is an important issue. Because RND volatility is a forecast of the volatility to be realized over an option's remaining lifetime we expect a high degree of autocorrelation in observations for a specific option contract on consecutive dates. In regressions with Realized volatility as the dependent, any future return shock that enters the residual for some date t will also be part of the residual for any observation for the same option on an earlier date. Indeed, we find first order autocorrelation in the OLS residuals above 0.9 in most cases. Beyond the obvious likely autocorrelation in the residuals for a given option contract, one should also anticipate strong correlation between the residuals for two options with different maturities that are observed on the same date. In short, the mixed time series cross section structure of the sample can be expected to induce correlation between the residuals from any two observations for options whose remaining lifetimes overlap. As is well-known, correlation in the residuals does not bias the coefficient point estimates, but the OLS standard errors are inconsistent.

Newey and West (1987) developed a procedure for correcting the standard errors from a regression that suffers from autocorrelation of a general form. It is related to methods proposed by White (1980) and Hansen (1982) for other failures of the OLS assumption of i.i.d. residuals. The essence of the approach is to model the covariance between the residuals from different dates as inversely related to the distance between them in time. I have adapted this approach to the more complicated correlation structure in my data set.

Let $\mathbf{X}$ be the matrix of the explanatory variables, with $\mathbf{X_i}$ as the row vector for observation i. The OLS coefficient vector is $\mathbf{b}$ and the vector of dependent variables is $\mathbf{y}$, with individual element y_i. The residual for observation i is $u_i \equiv y_i - \mathbf{X_i}'\mathbf{b}$.

Denote the true covariance matrix for the disturbance vector $\boldsymbol{\varepsilon}$ as $\boldsymbol{\Omega}$. The true covariance matrix for the estimated coefficients in a sample with N observations is given by

$$\text{Var}[\mathbf{b}] = N\,(\mathbf{X}'\mathbf{X})^{-1}\,(\mathbf{X}'\boldsymbol{\Omega}\mathbf{X})\,(\mathbf{X}'\mathbf{X})^{-1} \qquad (12)$$

Newey and West show that in a basic time series with one observation per date and an assumed autocorrelation structure such that residuals from observations separated by a lag of up to L periods are correlated but correlation is zero beyond that, a consistent estimator of **(X'ΩX)** can be obtained by

$$S_* \;=\; \frac{1}{N}\sum_i u_i^2\; X_i X'_i \;+\; \frac{1}{N}\sum_{j=1}^{L}\sum_{i=j+1}^{N} w_j\, u_i u_{i-j}(\,X_i X'_{i-j} +\; X_{i-j} X'_i) \qquad (13)$$

where the weight w_j is proportional to the number of periods of overlap as a fraction of L+1.[13] E.g., for two observations L periods apart, $w_j = 1/(L+1)$. This weighting structure produces a consistent estimator of the true matrix **X'ΩX** and constrains S_* to be positive semi-definite in finite samples. Newey and West also state that their definition for a weighting factor with these desired properties is not unique.

I have modified the Newey-West approach in (13) to adapt it to the more complicated correlation structure in my data set. The double summation adds a term to the error covariance matrix for every pair of observations for which correlation is expected. For this problem, that will be every pair of observations for which the remaining lifetimes of the two sets of options overlap.

Next, the weighting factor needs to be redefined appropriately to capture the assumption that the covariance between two residuals is proportional to the fractions of their lifetimes that overlap. The dependent variables are based on the average squared return over the option contract's remaining lifetime, $D_i \equiv (exp_i - date_i)$, where $date_i$ is the date of the observation and exp_i is the expiration date of the options in the ith observation. A return shock on a single future date contributes to the variance of y_i with a weight of $1/D_i$. The effect of that shock will also enter the variance of a different observation's y_j if it occurs in the future lifetime of the observation j options, with weight $1/D_j$. Thus the weighting factor w_j for this structure is given by

$$w_j \;=\; max\left(\frac{\min(exp_i\,,exp_j) - \max\,(date_i\,,date_j)}{(D_i D_j)^{1/2}}, 0\right) \qquad (14)$$

The definition of S_* for our regression sample is

$$S_* \;=\; \frac{1}{N}\sum_i u_i^2\; X_i X'_i \;+\; \frac{1}{N}\sum_{j=1}^{N-1}\sum_{i=j+1}^{N} w_j\, u_i u_{i-j}(\,X_i X'_{i-j} +\; X_{i-j} X'_i) \qquad (15)$$

and the coefficient standard errors are computed using

$$\text{Est. Var}[\mathbf{b}] \;=\; N\,(\mathbf{X'X})^{-1}\,\mathbf{S_*}\,(\mathbf{X'X})^{-1} \qquad (16)$$

Regressions with All Variables

Table 4 presents the results from four regressions run with the full set of explanatory variables, with t-statistics corrected using the method just described. The first two columns are the

[13] See Greene (2012, p.486).

coefficients and t-statistics with RND volatility and Realized volatility as the dependent variables and the full set of regressors. The rightmost columns present the results from the same regressions but adding yesterday's volatility premium as an explanatory variable. We do not favor this specification for the RND volatility regression because, while yesterday's volatility premium is obviously highly significant, this difference is a large part of what we are trying to explain with the exogenous factors. We include it here simply to show the effect on the other variables of adding this variable. The insignificant coefficient on the volatility risk premium in the Realized volatility regression is not surprising, although its negative sign is a little bothersome. What this variable is meant to measure is the outcome of the risk neutralization process, which is not a prediction of future volatility at all.

The table features some key results and some interesting differences between RND volatility and Realized volatility. Among the key results, one of the strongest, here and in other specifications we explored, is that the volatility forecast from the GARCH model is consistently highly significant, both for RND volatility and Realized volatility. GARCH appears to provide an important element in the market's prediction of volatility under the risk neutral measure. However, the coefficient is more than twice as large for Realized volatility as for RND volatility and both are well below 1.0.

Trading over a wider range on Date t is associated with higher risk neutral volatility and even more strongly with subsequent realized volatility. Whether the date t return fell in either the left or right tail had no significant effect on subsequent realized volatility. The t-statistics are much higher in the RND volatility regression, but the sign is wrong for a left tail event. A negative tail return multiplied by a positive coefficient implies that a bad shock to the downside lowered RND volatility.

For RND volatility, all of the Historical returns variables except the average left tail return came in with the expected signs and were significant, most highly so. By contrast, only three of the nine had significant explanatory power for Realized volatility.

Among the variables related to risk neutralization, strength in the Michigan Survey, the Baker-Wurgler investor confidence measure and the stock market price/earnings multiple were associated with lower RND volatility, although only the Michigan Survey received a significant coefficient. The Michigan sentiment measure also was significant for Realized volatility, but Baker-Wurgler investor sentiment came out strongly significant with the wrong sign. The Baa-Aaa bond yield spread was not significant for RND volatility but it had the expected positive sign and was nearly significant for Realized volatility.

In contrast to the correlation results in Table 3, the two variables related to accuracy of the GARCH model were insignificant, and 3 of 4 had the wrong signs.

Finally, comparing the results for RND volatility without and with last period's volatility premium as a regressor, the general impression is that the most significant variables in the left two columns without this variable remained so when it was added. There were changes in coefficient sizes and significance levels, but only a couple of variables changed sign and only the Date t right tail return went from significantly positive to significantly negative.

Including Historical volatility along with GARCH in the same specification allows for the expected result that GARCH would dominate Historical volatility in the Realized volatility equation, since it should be a superior forecaster. Either variable or both could be important to investors who may have less ability to construct and apply sophisticated econometric models to volatility prediction than financial econometricians do. And, indeed, both GARCH and Historical volatility do get significantly positive coefficients in the RND equation, but with the GARCH coefficient being more than three times as large. However, in the Realized volatility equation, Historical volatility has a negative coefficient even though the simple correlation between the two is 0.857.

Section 5: How Forward Looking is RND Volatility?

Theoretically, option prices are a function of the expected volatility from date t through expiration, but market makers, and traders who plan to hold an option only until the market moves to some target value, will have shorter forecasting horizons. The last section examined forming those volatility expectations from objective information that was available at date t, either from published statistics or from analysis of past returns data. One often-cited advantage of implying out parameter values from option market prices is that investors can have important information regarding future events and economic conditions that goes beyond what is contained in the current and historical data record, and this can be accessed by analyzing prices in the market.

To explore to what extent RND volatility is a forward looking forecast rather than just aggregating past information, and over what horizon, realized volatility was regressed on each of the three main volatility measures, RND volatility, GARCH, and Historical volatility, separately and in combination, over periods ranging from 1 day to the full life of the option. Table 5 presents the results. The table also reports the coefficients on these three variables in regressions with all of the explanatory variables included.

The first line in Table 5 shows that in a univariate regression, RND volatility is highly significant at every horizon with a coefficient slightly larger than 1.0. R^2 is greater than 0.50 for fixed horizons of 1 week (5 trading days) and above. The risk neutral density clearly contains information about the actual volatility that will be realized in the future. Comparing the results for different forecast horizons, the best explanatory power appears to be not in the few days immediately following date t, or over the full period to expiration. The largest coefficients and R^2 statistics are in the 5 to 10 day regressions.

One of the strongest results in the table is that the GARCH model forecast is highly significant at all horizons and in every regression it appears in. Adding the GARCH forecast to the specification in the second regression reduces the coefficient and t-statistic on RND volatility. The GARCH model coefficients and t-statistics are larger than those for RND volatility, but both variables are highly significant, which means that, in addition to the market's risk preferences, RND volatility does contain valuable information about future realized volatility that is not captured by GARCH..

By contrast, while Historical volatility is highly significant when run on its own, it becomes insignificant when combined with either RND volatility or GARCH, and with a negative coefficient in most cases.

With all three volatility variables in the regression, again GARCH gets the largest weight and statistical significance, but RND volatility also contributes significant explanatory power, while Historical volatility comes in significantly negative. The coefficients on RND volatility and GARCH are approximately equal over very short horizons of 1-3 days, but RND volatility does less well for longer horizons, while the performance of the GARCH forecast does not drop off, except when the forecast horizon is extended to the full period to option expiration.

Finally, the last four lines in Table 4 report the coefficients on the three volatility variables in regressions with the full set of explanatory variables (not including the previous day's risk premium). The multivariate results are fully consistent with the first part of the table, although both GARCH and RND volatility have smaller coefficients and lower t-statistics than in regressions without the extra variables.

Table 5 provides evidence that RND volatility contains useful information about future realized volatility beyond what can be captured by a GARCH model, and both are significantly better forecasts than vanilla historical volatility.

If GARCH is actually about the best one can do in forecasting future realized volatility using historical returns and RND volatility fully impounds GARCH, then the market is processing the available historical information correctly. And if RND volatility also impounds relevant information about the future that is not contained in historical returns, such as knowledge of an impending major election that can be expected to create increased volatility, it could be a better predictor than GARCH. But RND volatility also reflects variations in risk attitudes, so even if it does impound the best possible volatility prediction, it is a noisy forecast of future realized volatility. In that case, one expects RND volatility to be a strong predictor on its own, and to contribute significant explanatory power in combination with GARCH, but not to drive out GARCH when the two are run together. That is exactly what we have found.

Section 6: Subsample Results

This section explores the robustness of the relationships uncovered in the last section by breaking down the sample according to time to maturity and into subperiods that cover quite different economic environments.

Table 6 runs the regression of RND volatility on the full set of explanatory variables separately for short maturity contracts with less than 75 days to expiration and for longer maturity contracts. The strong results from the full sample regression largely hold in both maturity subsamples. The only coefficient to change sign, for example, is on the average size of left tail events in the last two years, which is insignificant in both subsamples. The average size of a right tail event becomes significant in the longer maturity regression, while the coefficient on GARCH RMSE becomes insignificant (with the wrong sign in both regressions). The results for short term and longer term options are, in fact, remarkably similar.

Table 7 compares regressions for both RND volatility and Realized volatility on all variables over three different time periods. The first is from January 1996 through June 2003, an interval that included the Russian debt crisis and the Internet bubble followed by the bear market of 2001-3. It was a period of medium volatility with a sharp run-up in stock prices in the first half and a sharp drop in the second half. The second subsample runs from July 2003 through December 2007. This was a time of persistently rising stock prices with extraordinarily low volatility. The third subperiod from January 2008 to April 2011 includes the market crash in the fall of 2008, when volatility soared to extreme heights, followed by the aftermath.

Not surprisingly, there is dispersion across subperiods in the results from both the RND volatility and Realized volatility equations, although the important results from the full sample regressions largely hold in the subperiods. The differences tend to be changes in statistical significance levels rather than sharp changes in the signs of the fitted coefficients.

The coefficient on the GARCH forecast remains positive and highly significant for both volatility measures, except for Realized volatility in the second subperiod. The coefficient on Historical volatility was positive and significant in all three RND volatility regressions, but insignificant in two of them for Realized volatility and significant with the wrong sign in the last one.

The Date t and historical return variables were negative in nearly all cases, and significant except for Date t return and Realized volatility in the last two subperiods. These variables were consistently more significant in the RND volatility than the Realized volatility equations.

The Date t trading range minus the absolute return was positive in all six regressions and highly significant in five of them. A large average daily range in the past increased RND volatility significantly in every subperiod, but Realized volatility only in the earliest one. The range covered by the index over the previous 25 days did not perform consistently. The coefficients in the RND volatility regressions were positive, and significant overall, but only significant in the final subperiod, while for Realized volatility, this variable was significantly positive in the earliest subperiod and significantly negative in the last.

For the most part, the left-tail related variables did not perform well. On Date t, a left tail return received a significantly positive coefficient in all three RND volatility regressions, meaning that an unusually large negative return was associated with a decrease in risk neutral volatility. More frequent left tail events in the previous 2 years increased both risk neutral and realized volatility overall, but results varied widely across subperiods. Similarly, the average size of left tail events showed inconsistent coefficient estimates and more than half had the wrong (i.e., positive) sign. Right tail events showed a similar lack of consistency between RND volatility and Realized volatility equations and over time.

Among the Risk neutralization variables, a higher value for the Michigan survey consumer sentiment variable reduced both RND volatility and Realized volatility significantly, with a couple of exceptions, while the Baker-Wurgler measure of investor confidence was more ambiguous, with both volatility measures increasing when this confidence measure was higher in the middle subperiod. By contrast, the P/E ratio for the S&P index had the right (negative) sign for both RND and Realized volatility in Table 4, but neither was significant. Here, that variable has consistently negative coefficients in all of the regressions, and half of them are significant. Overall, the coefficient on the bond yield spread was insignificant for RND volatility but positive and significant for Realized volatility, while in the subsample breakdown in Table 7, the coefficient is significantly positive for Realized volatility only in the final subperiod, which is consistent with default risk becoming a more important factor after 2008.

Finally, coefficients on the GARCH model errors and the GARCH RMSE in the subsample regressions continue to be largely insignificant and anomalous as they were in Table 4. A negative coefficient on GARCH RMSE implies that RND volatility is lower when the GARCH model has been less accurate in the recent past, and a negative coefficient on the GARCH error means that RND volatility was lower still when GARCH underpredicted volatility.

To summarize the results of these subsample regressions, there was not much difference between short and long maturity contracts in their sensitivities to the explanatory variables. More diverse results showed up when the sample was broken into subsamples by time period. High predicted volatility from both the GARCH and historical volatility models were associated with significantly higher risk neutral and realized volatilities in most cases, although the relationship was less strong for historical volatility in the Realized volatility regressions. The return variables showed the expected strongly negative correlation between return and volatility, while absolute return was more ambiguous. The variables relating to the trading range, intraday and over a longer historical period, were for the most part strongly positive: trading over a wider range increased RND volatility, and Realized volatility too, in most cases. The variables measuring frequency and size of tail events showed inconsistent results, with some coefficients significant, but with possibly different signs in different subperiods, and quite a few anomalous values. The results for the Risk neutralization variables that were hypothesized as proxies for the market's tolerance for risk bearing were promising, though not overpoweringly strong. In most cases, more positive sentiment seems to be associated with lower volatility. However, the Aaa-Baa bond yield spread and the two variables measuring the accuracy of the GARCH model did not add much explanatory power.

Section 7: Concluding Comments

We set out to develop stylized facts about what things influence risk neutral volatility, including factors related to how investors might form estimates of the true probability distribution governing future returns and factors related to the risk neutralization process that transforms the expected true "P" distribution into the risk neutral "Q" distribution which determines market prices. Several broad questions were identified for exploration. Let us summarize what we have uncovered.

The first question was what return and volatility-related factors are most important to investors in forming forecasts of the empirical probability density that are embedded in the RND? Under Black-Scholes assumptions, this amounts to getting the most accurate possible estimate of the instantaneous volatility of the underlying lognormal diffusion process assumed for returns. The natural choice of estimator for a constant volatility diffusion is simple historical volatility. But given that observed returns processes differ significantly from lognormal diffusions, including strong evidence that volatility varies over time, investors might well take other approaches to estimating volatility and also take other volatility-related aspects of returns behavior, such as trading ranges and tail events, into consideration separately.

Historical volatility is the plain vanilla way to estimate volatility, but a GARCH model uses the same returns information in a more sophisticated way so, as was shown in Table 5, it provides more accurate predictions of future volatility. It is not surprising, then, that the GARCH model forecast was highly significant in regressions for Realized volatility, while historical volatility was not, and was even estimated with a negative coefficient. By contrast, while the GARCH forecast was also highly significant in the RND volatility regressions, so was historical volatility. Investors appear to pay attention to both in pricing options.

A number of volatility-related factors that investors could reasonably want to take account of involve analyzing historical returns. These include historical volatility, recent returns and trading ranges, tail events and the past accuracy of a GARCH model, but how far back investors look is an open question. To shed light on this issue and to decide on the best specifications for our explanatory variables, we analyzed t-statistics and R^2s from regressions of RND volatility on each variable separately, and in combination with the other exogenous variables. We used a range of sample lengths, from 2 days to 2 years, and also a range of cutoff values for defining tail events. Although this exercise is inherently somewhat subjective, the results were quite illuminating in most cases. The information contained in past returns and absolute returns was strong at all lags, but the greatest contribution to RND volatility seems to be from returns over the most recent few days. Historical volatility was uniformly significant in the univariate regressions, but only at longer horizons when run with other variables (that included GARCH). Variables measuring the average trading range within a day and the full range traversed over a span of time both seemed to do best in samples of medium length, leading us to define the average daily range over a period equal to the remaining lifetime of each specific option, and the range of the sample path over the previous 25 trading days. The two measures of GARCH model accuracy also seemed to do best for sample periods equal to the option's remaining life.

Results were mixed for the left and right tail variables, both in the exploratory regressions to settle on tail definitions, and then in the subsequent regressions. There were anomalous signs, for example, results suggesting that worse recent losses in the left tail reduced RND volatility, and major differences between univariate and multivariate regression results and across sample lengths. It did appear that longer historical samples performed better than shorter ones and that setting the tail cutoffs at 0.5% was too restrictive. In the end, the results for the tail variables in Tables 4, 6 and 7 might be called "tantalizing" because of high statistical significance in many cases, but not consistent enough to be called "definitive."

Among the strongest results is that the well-known negative relationship between return and volatility was clearly manifested here. Both the Date t return and the return over the previous week were estimated with highly significant negative coefficients overall and, with a single exception, in every subperiod regression. However, the coefficients were larger and more significant in the RND volatility regressions than in those for Realized volatility. These differences suggest that empirical ("P measure") volatility is correlated with returns, but so is the risk neutralization process that transforms the P measure into the Q measure, which increases the effect.

Results for the variables connected to the trading range were somewhat similar to those with returns. Trading over a wide range that produced little change in the index level on Date t was highly significant in increasing both RND and Realized volatility. But a large daily range on average and a large total range in the recent past also increased RND volatility substantially, while having no significant effect on realized volatility. This suggests that these variables may be more strongly connected to investors' and market makers' risk attitudes than to their objective forecasts of future volatility.

Along with the explanatory variables derived from current and past returns, we considered six that were specifically chosen for their possible correlation with risk tolerance. RND volatility was negatively correlated with the Michigan survey of consumer sentiment and the Baker-Wurgler index of investor sentiment, with the former showing strong performance, both for RND volatility and Realized volatility. It would hardly be surprising to find consumers and investors less confident in times of high market volatility. However, the effect appears distinctly stronger for RND volatility than Realized volatility, suggesting that there is an additional impact of sentiment on risk neutralization. The price/earnings ratio of the S&P 500 and the credit spread in the bond market are direct measures of the market's willingness to bear risk. The coefficient on the P/E ratio was negative for both volatility measures overall and in every subperiod, which suggests that this variable may be a useful reflection of market sentiment, although it was only significant in a few cases. The bond yield spread results were odd overall and inconsistent in the breakdown across subperiods. Finally, the two variables designed to measure investors' confidence in the GARCH model's predictions of future volatility based on its accuracy in the past did not add explanatory power. The coefficients were mostly insignificant with anomalous signs.

In the end, we have seen that RND volatility and Realized volatility are strongly but not perfectly correlated (rho $\cong$ 0.6). Forecasts from a GARCH model, that only tries to predict future realized volatility, are highly significant in explaining RND volatility, but adding variables to capture the possible separate influence of other aspects of returns that investors may care about, such as the past trading range, and factors specifically targeted toward the risk neutralization process, greatly increases explanatory power. Beyond expected future volatility, these factors are significant determinants of how options are priced in the real world.

References

Baker, M. and J. Wurgler (2006). "Investor Sentiment and the Cross-Section of Stock Returns." Journal of Finance 61, 1645-1680.

Birru, J. and S. Figlewski (2012). "Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008." Journal of Financial Markets 15, 151-180.

Bliss, R. and N. Panigirtzoglou (2002). "Testing the Stability of Implied Density Functions." Journal of Banking and Finance 26, 381-422.

Bliss, R. and N. Panigirtzoglou (2004). "Option Implied Risk Aversion Estimates." Journal of Finance 59, 407-446.

Bollerslev, T. and V. Todorov (2011). "Tails, Fears and Risk Premia." Forthcoming, Journal of Finance.

Breeden, D. and R. Litzenberger (1978). "Prices of state-contingent claims implicit in option prices." Journal of Business 51, 621–652.

Carr, P. and L. Wu (2008). "Variance Risk Premiums." Review of Financial Studies 22, 1311-41.

Chicago Board Options Exchange (2003). VIX CBOE Volatility Index.

Christoffersen, P., K. Jacobs and B. Y. Chang (2011). "Forecasting with Option Implied Information." Working paper, available at http://ssrn.com/abstract=1969863.

Figlewski, S. (2010). "Estimating the implied risk neutral density for the U.S. market portfolio." In: Bollerslev, T., Russell, J.R., Watson, M. (Eds.),Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle. Oxford University Press, Oxford, UK.

Garman, M.B. and M.J. Klass (1980). "On the estimation of security price volatilities from historical data." Journal of Business 53, 67–78.

Gatheral, J. (2006). The Volatility Surface: A Practitioner's Guide. Wiley Finance.

Giamouridis, D. and G. Skiadopoulos (2012). "The Informational Content of Financial Options for Quantitative Asset Management." In: Scherer, B. and Winston, K. (Eds.), The Oxford Handbook of Quantitative Asset Management. Oxford University Press, Oxford, U.K.

Glosten, L.R., R. Jagannathan, and D.E. Runkle (1993). "On the relation between the expected value and the volatility of the nominal excess return on stocks." Journal of Finance 48, 1779–1801.

Greene, W. (2012). Econometric Analysis, 7th ed. Prentice-Hall, Upper Saddle River, NJ.

Hansen, L. P. (1982). "Large Sample Properties of Generalized Method of Moments Estimators." Econometrica 50, 1029-1054.

Harrison, J. M. and D.M. Kreps (1979). "Martingales and arbitrage in multiperiod securities markets." Journal of Economic Theory 20, 381-408.

Jackwerth, J.C. (2004). Option-Implied Risk-Neutral Distributions and Risk Aversion. Charlotteville, Research Foundation of AIMR.

Jiang, G. and Y. S. Tian (2007). "Extracting Model-Free Volatility from Option Prices: An Examination of the VIX Index." Journal of Derivatives 14, 35-60.

Kozhan, R., A. Neuberger and P. Schneider (2010). "Understanding Risk Premia in Option Prices." Working paper, available at http://ssrn.com/abstract=1571700.

Parkinson, M. (1980). "The extreme value method for estimating the variance of the rate of return." Journal of Business 53, 61-68.

Poon, S.-H. and C.W.J. Granger (2003). " Forecasting Volatility in Financial Markets: A Review." Journal of Economic Literature 41, 478-539.

Ross, S. (2011). "The Recovery Theorem." NBER Working Paper, available at http://www.nber.org/papers/w17323.

White, H. (1980). "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." Econometrica 48, 817–838.

Wurgler, J. "Investor sentiment data (annual and monthly), 1934-2010." available at http://people.stern.nyu.edu/jwurgler/ .

Zhou, H. (2009). "Variance Risk Premia, Asset Predictability Puzzles, and Macroeconomic Uncertainty." Federal Reserve Board Working Paper, available at http://ssrn.com/abstract=1400049.

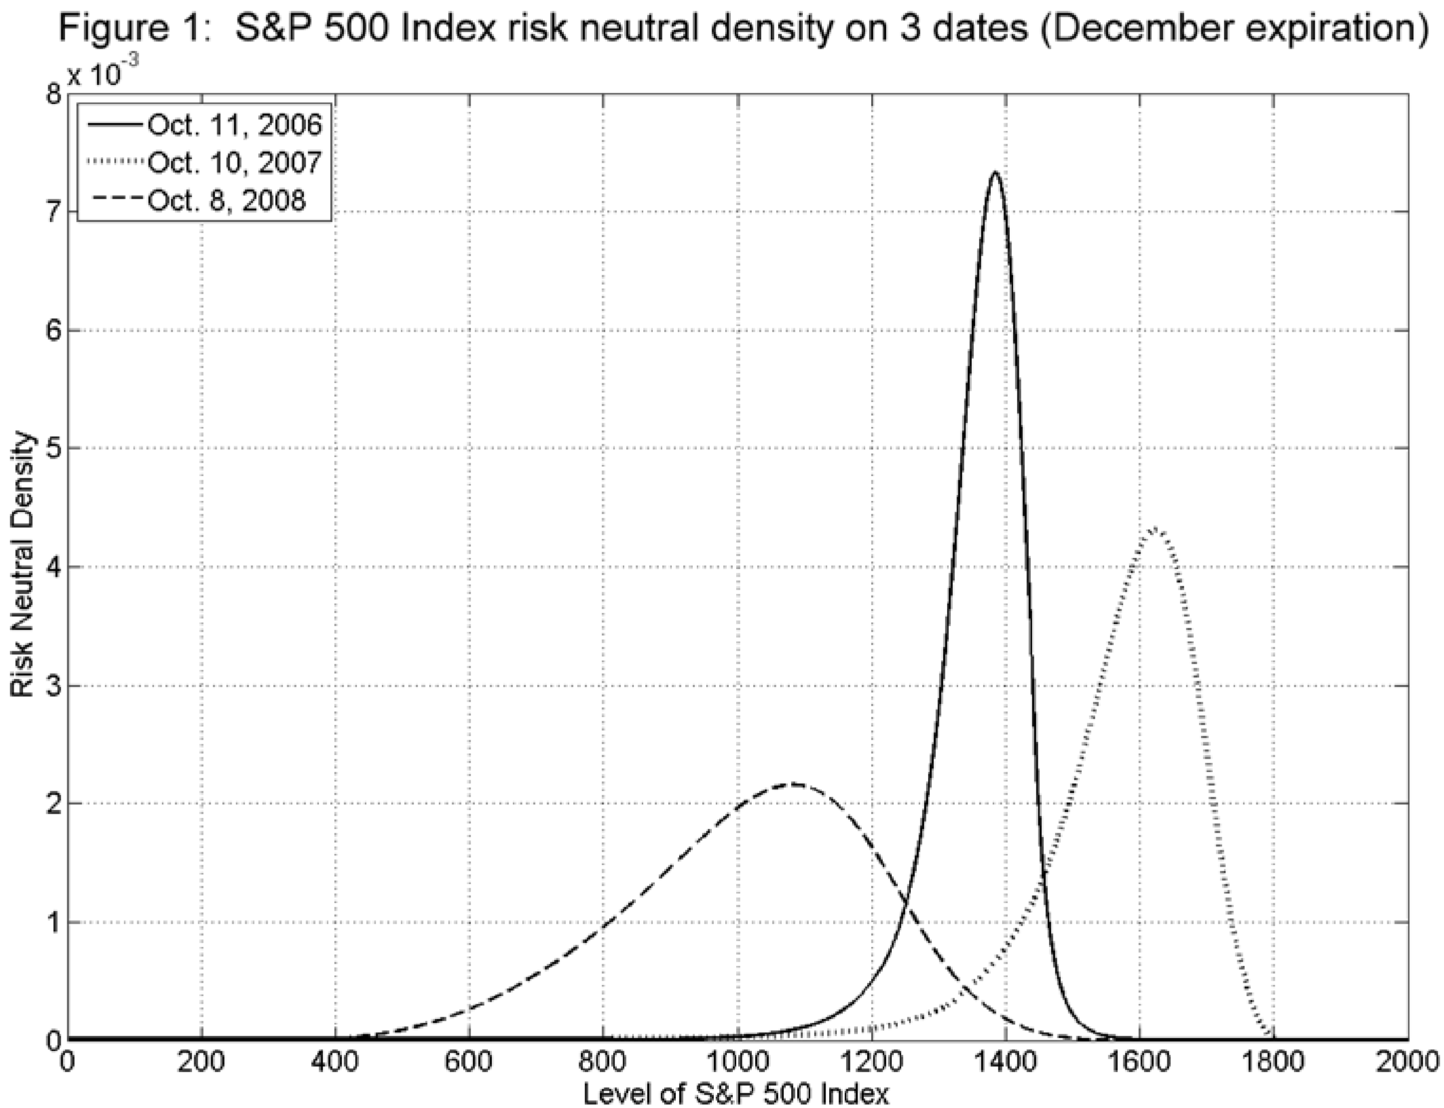
Figure 1: S&P 500 Index risk neutral density on 3 dates (December expiration)
x 10^-3
Oct. 11, 2006
Oct. 10, 2007
Oct. 8, 2008
Risk Neutral Density
Level of S&P 500 Index
0
1
2
3
4
5
6
7
8
0
200
400
600
800
1000
1200
1400
1600
1800
2000

Table 1. Alternative Tail Cutoff Definitions

Panel A Univariate Regressions

	date t	2 days	3 days	5 days	10 days	25 days	65 days	125 days	250 days	500 days	option life
Tail = 0.5%											
# of past left tail events	--	0.201	-0.038	-0.135	-0.590	-1.112	-1.370	-1.705	-0.944	1.787	-1.635
Size of past left tail events	-0.742	-0.794	-0.543	-0.460	0.034	0.605	1.011	1.911	0.402	-2.258	1.532
# of past right tail events	--	1.264	1.246	1.259	1.453	1.582	1.608	1.658	3.491	3.473	1.822
Size of past right tail events	1.896	2.099	2.160	2.151	2.267	2.487	2.864	2.731	3.981	5.495	2.914
Tail = 1.0%											
# of past left tail events	--	1.800	1.893	1.961	1.596	1.878	2.081	2.240	3.579	4.136	1.960
Size of past left tail events	-2.800	-3.097	-3.057	-3.052	-2.568	-2.830	-3.141	-4.098	-4.725	-5.965	-3.464
# of past right tail events	--	1.407	1.797	1.954	2.392	2.189	1.832	2.087	2.762	5.561	2.120
Size of past right tail events	2.216	2.436	3.421	3.664	4.197	4.762	5.971	8.446	4.823	3.642	6.344
Tail = 2.0%											
# of past left tail events	--	0.892	0.887	0.881	0.854	1.405	2.071	3.099	5.915	6.307	2.880
Size of past left tail events	-2.221	-2.040	-1.998	-1.999	-1.954	-2.372	-2.570	-3.322	-5.051	-6.099	-3.219
# of past right tail events	--	1.910	2.028	2.118	2.247	2.258	2.135	2.646	2.795	4.901	2.001
Size of past right tail events	3.873	4.232	4.375	4.653	4.798	5.013	7.477	9.328	4.445	4.339	7.402
Tail = 5.0%											
# of past left tail events	--	4.462	4.419	4.674	4.741	5.043	5.157	5.723	7.177	7.395	4.449
Size of past left tail events	-6.738	-7.415	-7.514	-8.025	-8.172	-10.401	-12.188	-9.250	-5.201	-5.164	-13.072
# of past right tail events	--	0.566	0.734	0.737	1.073	0.929	0.572	0.551	1.332	-1.161	-0.032
Size of past right tail events	3.987	4.351	4.700	4.950	5.558	6.280	7.829	7.380	6.031	6.220	9.447

Table 1. Alternative Tail Cutoff Definitions, continued

Panel B Multivariate Regressions

	date t	2 days	3 days	5 days	10 days	25 days	65 days	125 days	250 days	500 days	option life
Tail = 0.5%											
# of past left tail events	--	-2.802	-2.999	-3.178	-2.939	-3.696	-4.434	-4.237	-0.720	-3.414	-4.510
Size of past left tail events	2.598	5.231	5.144	4.660	3.917	4.019	4.051	3.386	-0.512	-5.014	3.464
# of past right tail events	--	-0.348	-0.169	-0.059	0.907	1.115	0.274	-1.659	-0.821	-2.438	-0.065
Size of past right tail events	2.327	0.168	0.393	0.526	1.612	1.761	0.919	-0.240	0.787	5.980	1.244
Tail = 1.0%											
# of past left tail events	--	-1.737	-2.312	-2.951	-2.815	-1.099	-1.636	-1.754	1.218	1.340	-1.721
Size of past left tail events	1.230	4.012	4.216	4.880	5.013	3.392	4.163	3.832	-1.093	-1.710	3.189
# of past right tail events	--	2.665	2.739	3.213	3.559	2.845	2.249	1.189	0.384	2.735	1.948
Size of past right tail events	4.209	0.857	0.310	0.231	1.163	2.147	0.980	0.497	-0.586	0.391	2.009
Tail = 2.0%											
# of past left tail events	--	-2.885	-3.349	-3.640	-2.023	-0.302	0.836	1.890	6.534	7.135	1.692
Size of past left tail events	3.277	5.839	6.734	7.989	5.811	3.844	4.646	3.607	-0.226	0.157	2.626
# of past right tail events	--	0.989	1.025	0.921	2.292	2.632	0.932	-0.063	-0.897	1.627	1.132
Size of past right tail events	3.146	-0.066	-0.248	-0.867	0.449	2.456	0.645	1.310	1.410	2.324	2.285
Tail = 5.0%											
# of past left tail events	--	-0.151	-1.389	-2.287	-0.065	1.874	2.561	3.835	6.371	5.281	3.700
Size of past left tail events	1.084	2.948	4.484	4.392	1.681	0.782	1.846	-1.728	0.263	0.996	-1.321
# of past right tail events	--	-0.705	-0.225	0.026	1.223	1.287	0.186	-1.024	-0.114	0.332	1.056
Size of past right tail events	3.546	-0.551	-0.234	0.220	1.675	3.817	1.879	2.609	1.020	1.281	3.667

Notes: The Table reports t-statistics on four variables measuring returns falling into the left and right tails of the returns distribution from regressions on RND volatility. Results for univariate regressions are shown in Panel A, while Panel B reports on regressions that include all of the explanatory variables in the study. The first column considers only the single current date t, while the other columns evaluate tails outcomes over historical periods of different lengths. The rightmost column sets the historical sample length equal to the remaining lifetime of the options used to compute the RND for that observation. The t-statistics are corrected for the cross correlation caused by overlapping option lives, as described in the text.

Table 2. Alternative Lag Definitions for Variables Computed from Historical Returns

Panel A Regression t-statistics

Variable		2 days	3 days	5 days	10 days	25 days	65 days	125 days	250 days	500 days	option life
Return t-n to t-1	**univariate**	-3.171	-2.996	-2.877	-2.516	-2.722	-3.742	-4.760	-5.362	-4.735	-4.225
	multivariate	-9.177	-11.693	-10.674	-4.881	-2.221	-3.435	-4.076	-2.167	0.311	-4.508
Absolute return t-n to t-1	**univariate**	7.699	7.172	8.631	7.588	7.748	6.499	5.592	5.206	4.224	5.500
	multivariate	2.779	2.369	3.567	3.401	2.172	-0.964	-2.341	0.052	3.133	-0.200
Historical volatility t-n to t-1	**univariate**	9.524	11.376	13.581	15.011	16.990	14.603	10.983	5.022	5.033	20.404
	multivariate	0.076	-0.630	-2.806	-0.219	3.078	4.134	4.632	1.624	1.763	5.976
Avg daily range t-n to t-1	**univariate**	12.453	13.912	14.248	14.202	17.095	15.113	10.644	5.163	5.718	22.170
	multivariate	4.550	4.227	4.406	3.706	2.964	3.384	4.205	0.937	2.041	5.050
Range over t-n to t-1	**univariate**	11.414	11.233	11.487	12.684	14.677	15.819	11.431	5.950	5.284	6.505
	multivariate	2.372	0.858	-0.159	3.548	3.459	2.017	3.724	1.972	3.275	3.138
GARCH error t-n to t-1	**univariate**	0.552	0.329	0.485	0.901	1.268	2.423	7.437	5.928	5.379	4.452
	multivariate	0.565	0.082	-1.433	-1.370	-0.738	-2.243	-2.150	-0.173	-1.876	-1.672
GARCH RMSE t-n to t-1	**univariate**	9.547	9.698	8.206	7.150	7.506	9.268	9.627	4.571	4.404	16.725
	multivariate	-1.879	-2.278	-4.920	-5.362	-1.579	-4.052	1.661	1.258	1.516	1.345

Table 2, cont. Alternative Lag Definitions for Variables Computed from Historical Returns

Panel B Regression R-squared Statistics

Variable		2 days	3 days	5 days	10 days	25 days	65 days	125 days	250 days	500 days	option life
Return t-n to t-1	univariate	0.010	0.012	0.018	0.035	0.107	0.279	0.387	0.399	0.311	0.238
	multivariate	0.935	0.937	0.938	0.936	0.935	0.935	0.937	0.934	0.933	0.935
Absolute return t-n to t-1	univariate	0.259	0.276	0.242	0.231	0.266	0.352	0.380	0.411	0.236	0.289
	multivariate	0.937	0.937	0.938	0.938	0.938	0.937	0.938	0.937	0.940	0.937
Historical volatility t-n to t-1	univariate	0.394	0.492	0.601	0.689	0.766	0.789	0.742	0.507	0.309	0.778
	multivariate	0.934	0.934	0.935	0.934	0.937	0.938	0.941	0.934	0.935	0.940
Avg daily range t-n to t-1	univariate	0.574	0.622	0.668	0.719	0.770	0.767	0.688	0.460	0.250	0.756
	multivariate	0.937	0.937	0.938	0.939	0.939	0.938	0.941	0.936	0.937	0.939
Range over t-n to t-1	univariate	0.553	0.572	0.591	0.622	0.668	0.694	0.669	0.483	0.319	0.426
	multivariate	0.935	0.935	0.935	0.936	0.937	0.936	0.938	0.936	0.937	0.936
GARCH error t-n to t-1	univariate	0.001	0.000	0.002	0.012	0.052	0.186	0.514	0.578	0.450	0.166
	multivariate	0.938	0.938	0.938	0.938	0.938	0.938	0.938	0.938	0.938	0.938
GARCH RMSE t-n to t-1	univariate	0.372	0.418	0.459	0.532	0.611	0.673	0.684	0.486	0.331	0.671
	multivariate	0.938	0.938	0.940	0.940	0.938	0.941	0.939	0.938	0.939	0.938

Notes: The Table reports t-statistics and R^2 results from regressions of RND volatility on each of seven variables that are constructed from the series of returns from date t-n to t-1for different values of n. The rightmost column sets the historical sample length equal to the remaining lifetime of the options used to compute the RND for that observation. For each variable the first line shows results from univariate regressions and the second from multivariate regressions that include all of the explanatory variables in the study. t-statistics are shown in Panel A and R^2 in Panel B. The t-statistics are corrected for the cross correlation caused by overlapping option lives, as described in the text.

Table 3. Sample Correlations

	RND volatility	Realized volatility to expiration	Date t return	Date t absolute return	GARCH to expiration	Date t trading range	Date t range minus absolute return	Date t left 2% tail return	Date t right 2% tail return	Last week return	Last week absolute return	Historical volatility (65 days)
Dependent variables												
RND volatility	1.000	0.613	0.976	-0.063	0.488	0.877	0.687	0.351	-0.031	0.068	-0.136	0.492
Realized volatility to expiration	0.613	1.000	0.639	-0.069	0.385	0.678	0.538	0.271	-0.055	0.052	-0.159	0.333
Date t return variables												
Date t return	-0.063	-0.069	1.000	-0.018	-0.035	-0.033	-0.025	0.251	0.358	0.411	-0.074	0.014
Date t absolute return	0.488	0.385	-0.018	1.000	0.571	0.805	-0.261	-0.261	0.384	-0.152	0.336	0.436
GARCH to expiration	0.877	0.678	-0.035	0.571	1.000	0.748	0.319	-0.059	0.102	-0.127	0.500	0.857
Date t trading range	0.687	0.538	-0.033	0.805	0.748	1.000	0.362	-0.205	0.293	-0.240	0.431	0.604
Date t range minus absolute return	0.351	0.271	-0.025	-0.261	0.319	0.362	1.000	0.078	-0.127	-0.152	0.173	0.298
Date t left 2% tail return	-0.031	-0.055	0.251	-0.261	-0.059	-0.205	0.078	1.000	0.010	0.130	-0.092	0.003
Date t right 2% tail return	0.068	0.052	0.358	0.384	0.102	0.293	-0.127	0.010	1.000	0.096	0.078	0.063
Historical returns variables												
Last week return	-0.136	-0.159	0.411	-0.152	-0.127	-0.240	-0.152	0.130	0.096	1.000	-0.033	0.002
Last week absolute return	0.492	0.333	-0.074	0.336	0.500	0.431	0.173	-0.092	0.078	-0.033	1.000	0.422
Historical volatility (65 days)	0.888	0.589	0.014	0.436	0.857	0.604	0.298	0.003	0.063	0.002	0.422	1.000
Avg daily range (past pd = option life)	0.870	0.564	0.026	0.402	0.809	0.569	0.294	0.014	0.067	0.028	0.409	0.920
Range over last 25 days	0.817	0.567	0.044	0.452	0.844	0.647	0.342	-0.001	0.115	0.041	0.466	0.787
# left 2% tail events last 2 years	0.517	0.291	-0.005	0.199	0.296	0.296	0.170	-0.073	0.009	-0.009	0.234	0.376
Avg left 2% tail return last 2 years	-0.579	-0.314	-0.014	-0.208	-0.374	-0.302	-0.164	0.009	-0.020	-0.006	-0.245	-0.464
# right 2% tail events last 2 years	0.380	0.169	0.048	0.133	0.226	0.186	0.094	0.015	0.054	0.080	0.147	0.334
Avg right 2% tail return last 2 years	0.485	0.130	0.034	0.130	0.281	0.180	0.088	0.010	0.025	0.052	0.162	0.378
Risk neutralization variables												
Michigan consumer sentiment	-0.444	-0.316	-0.015	-0.179	-0.401	-0.223	-0.083	-0.018	-0.043	-0.012	-0.168	-0.436
Baker-Wurgler investor sentiment	-0.138	0.076	-0.024	-0.003	-0.051	0.012	0.025	-0.015	0.023	-0.067	-0.002	-0.085
S&P 500 P/E ratio	-0.548	-0.379	-0.015	-0.237	-0.402	-0.345	-0.190	-0.007	-0.021	-0.020	-0.250	-0.566
BAA-AAA bond yield spread	0.762	0.535	0.023	0.369	0.750	0.498	0.230	0.027	0.061	0.011	0.343	0.850
GARCH error (past pd = option life)	0.407	0.371	0.012	0.343	0.561	0.425	0.153	-0.062	0.099	-0.018	0.232	0.365
GARCH RMSE (past pd = option life)	0.819	0.550	0.012	0.419	0.836	0.572	0.272	-0.012	0.070	-0.005	0.395	0.898
RND volatility premium date t-1	0.047	-0.255	0.042	-0.111	-0.412	-0.170	-0.100	-0.032	0.016	-0.001	-0.108	-0.115

Table 3. Sample Correlations, continued

	Avg daily range (past pd = option life)	Range over last 25 days	# left 2% tail events last 2 years	Avg left 2% tail return last 2 years	# right 2% tail events last 2 years	Avg right 2% tail return last 2 years	Michigan consumer sentiment	Baker-Wurgler investor sentiment	S&P 500 P/E ratio	BAA-AAA bond yield spread	GARCH error (past pd = option life)	GARCH RMSE (past pd = option life)	RND volatility premium date t-1
Dependent variables													
RND volatility	0.888	0.870	0.817	0.517	-0.579	0.380	0.485	-0.444	-0.138	-0.548	0.762	0.407	0.819
Realized volatility to expiration	0.589	0.564	0.567	0.291	-0.314	0.169	0.130	-0.316	0.076	-0.379	0.535	0.371	0.550
Date t return variables													
Date t return	0.026	0.044	-0.005	-0.014	0.048	0.034	-0.015	-0.024	-0.015	0.023	0.012	0.012	0.042
Date t absolute return	0.402	0.452	0.199	-0.208	0.133	0.130	-0.179	-0.003	-0.237	0.369	0.343	0.419	-0.111
GARCH to expiration	0.809	0.844	0.296	-0.374	0.226	0.281	-0.401	-0.051	-0.402	0.750	0.561	0.836	-0.412
Date t trading range	0.569	0.647	0.296	-0.302	0.186	0.180	-0.223	0.012	-0.345	0.498	0.425	0.572	-0.170
Date t range minus absolute return	0.294	0.342	0.170	-0.164	0.094	0.088	-0.083	0.025	-0.190	0.230	0.153	0.272	-0.100
Date t left 2% tail return	0.014	-0.001	-0.073	0.009	0.015	0.010	-0.018	-0.015	-0.007	0.027	-0.062	-0.012	-0.032
Date t right 2% tail return	0.067	0.115	0.009	-0.020	0.054	0.025	-0.043	0.023	-0.021	0.061	0.099	0.070	0.016
Historical returns variables													
Last week return	0.028	0.041	-0.009	-0.006	0.080	0.052	-0.012	-0.067	-0.020	0.011	-0.018	-0.005	-0.001
Last week absolute return	0.409	0.466	0.234	-0.245	0.147	0.162	-0.168	-0.002	-0.250	0.343	0.232	0.395	-0.108
Historical volatility (65 days)	0.920	0.787	0.376	-0.464	0.334	0.378	-0.436	-0.085	-0.566	0.850	0.365	0.898	-0.115
Avg daily range (past pd = option life)	1.000	0.762	0.387	-0.491	0.408	0.388	-0.398	-0.105	-0.635	0.801	0.412	0.921	-0.055
Range over last 25 days	0.762	1.000	0.346	-0.418	0.293	0.314	-0.341	-0.005	-0.440	0.654	0.466	0.731	-0.214
# left 2% tail events last 2 years	0.387	0.346	1.000	-0.562	0.173	0.287	0.203	0.016	-0.403	0.108	0.155	0.289	0.358
Avg left 2% tail return last 2 years	-0.491	-0.418	-0.562	1.000	-0.575	-0.723	0.227	0.000	0.539	-0.350	-0.117	-0.369	-0.316
# right 2% tail events last 2 years	0.408	0.293	0.173	-0.575	1.000	0.465	-0.219	-0.169	-0.558	0.317	0.088	0.244	0.248
Avg right 2% tail return last 2 years	0.388	0.314	0.287	-0.723	0.465	1.000	-0.533	-0.334	-0.205	0.415	0.055	0.320	0.340
Risk neutralization variables													
Michigan consumer sentiment	-0.398	-0.341	0.203	0.227	-0.219	-0.533	1.000	0.387	0.079	-0.671	-0.097	-0.430	-0.018
Baker-Wurgler investor sentiment	-0.105	-0.005	0.016	0.000	-0.169	-0.334	0.387	1.000	0.034	-0.327	0.011	-0.138	-0.161
S&P 500 P/E ratio	-0.635	-0.440	-0.403	0.539	-0.558	-0.205	0.079	0.034	1.000	-0.525	-0.087	-0.501	-0.189
BAA-AAA bond yield spread	0.801	0.654	0.108	-0.350	0.317	0.415	-0.671	-0.327	-0.525	1.000	0.256	0.829	-0.122
GARCH error (past pd = option life)	0.412	0.466	0.155	-0.117	0.088	0.055	-0.097	0.011	-0.087	0.256	1.000	0.512	-0.380
GARCH RMSE (past pd = option life)	0.921	0.731	0.289	-0.369	0.244	0.320	-0.430	-0.138	-0.501	0.829	0.512	1.000	-0.198
RND volatility premium date t-1	-0.055	-0.214	0.358	-0.316	0.248	0.340	-0.018	-0.161	-0.189	-0.122	-0.380	-0.198	1.000

Notes: The table shows the simple correlation coefficients among all of the variables considered in the study. The sample includes data from 10,152 date-option maturity pairs from January 4, 1996 through April 27, 2011, with bad data points excluded as described in the text.

Table 4: Full Sample Regressions

	RND volatility		Realized volatility		RND volatility		Realized volatility	
	coef	t-stat	coef	t-stat	coef	t-stat	coef	t-stat
Date t return variables								
Constant	10.496	5.078	21.856	1.833	1.918	4.155	27.199	1.997
Date t return	-0.190	-10.092	-0.171	-2.548	-0.273	-11.450	-0.124	-1.649
GARCH to expiration	0.324	8.455	0.802	3.840	0.847	49.315	0.479	2.600
Date t range minus absolute return	0.355	6.025	0.752	4.346	0.555	17.201	0.645	3.098
Date t left 2% tail return	0.444	4.009	-0.035	-0.090	1.048	10.508	-0.644	-1.136
Date t right 2% tail return	0.213	3.421	0.015	0.057	-0.321	-7.418	0.344	1.342
Historical returns variables								
Last week return	-0.005	-12.687	-0.004	-2.778	0.0001	0.475	-0.007	-3.116
Last week absolute return	0.003	4.332	-0.003	-0.819	0.0004	1.362	-0.001	-0.217
Historical volatility (65 days)	0.106	3.761	-0.202	-1.273	0.025	3.297	-0.152	-1.073
Avg daily range (past pd = option life)	1.894	4.690	1.483	0.993	0.306	3.404	2.263	1.230
Range over last 25 days	0.161	3.880	-0.184	-0.851	0.048	4.485	-0.108	-0.596
# left 2% tail events last 2 years	0.816	8.663	1.325	2.307	0.149	6.225	1.744	2.302
Avg left 2% tail return last 2 years	0.055	0.262	-0.781	-1.197	-0.023	-0.545	-0.836	-1.277
# right 2% tail events last 2 years	0.113	1.690	0.069	0.269	0.022	1.676	0.094	0.367
Avg right 2% tail return last 2 years	0.343	2.051	-2.091	-2.372	0.030	0.832	-1.990	-2.470
Risk neutralization variables								
Michigan consumer sentiment	-0.069	-3.996	-0.221	-1.849	-0.014	-4.035	-0.254	-1.952
Baker-Wurgler investor sentiment	-0.297	-1.218	3.163	2.685	-0.056	-1.164	2.984	2.698
S&P 500 P/E ratio	-0.120	-1.100	-0.430	-0.736	0.020	0.958	-0.473	-0.832
BAA-AAA bond yield spread	-0.189	-0.268	5.041	1.542	0.012	0.086	4.876	1.575
GARCH error (past pd = option life)	-0.001	-1.292	0.005	1.373	-0.0001	-0.763	0.004	1.296
GARCH RMSE (past pd = option life)	-0.0004	-1.358	-0.002	-1.270	0.0000	-0.363	-0.002	-1.255
RND volatility premium date t-1					0.813	45.384	-0.477	-1.521
R-squared	0.939		0.564		0.985		0.569	
NOBS	10152		10152		9395		9395	

Notes: Results from OLS regressions of RND volatility and Realized volatility from observation date through option expiration, regressed on the full set of explanatory variables, with t-statistics corrected for cross correlation caused by overlapping option lives, as described in the text.

Table 5. Regressions of Realized Volatility over Various Horizons on RND, GARCH, and Historical Volatility

	1-Day		2-Day		3-Day		5-Day		7-Day		10-Day		15-Day		To expiration	
	coef	t-stat	coef	t-stat	coef	t-stat	coef	t-stat	coef	t-stat	coef	t-stat	coef	t-stat	coef	t-stat
Coefficients from Regressions on Volatility Variables Only																
RND volatility date t	1.011	11.275	1.102	12.659	1.148	12.398	1.176	11.395	1.164	11.670	1.143	11.514	1.111	11.203	0.897	9.612
R-squared	0.218		0.353		0.439		0.522		0.548		0.555		0.535		0.376	
RND volatility date t	0.441	6.457	0.484	5.712	0.464	5.359	0.370	3.257	0.346	2.729	0.328	2.377	0.286	1.836	0.112	0.571
GARCH to expiration	0.590	6.819	0.639	5.744	0.707	6.038	0.833	5.483	0.847	5.142	0.843	4.716	0.854	4.317	0.812	3.665
R-squared	0.238		0.386		0.485		0.595		0.629		0.639		0.623		0.462	
RND volatility date t	1.079	7.324	1.176	7.663	1.194	7.307	1.184	6.948	1.130	6.987	1.068	6.777	0.959	6.570	0.621	4.453
Historical volatility	-0.050	-0.629	-0.055	-0.659	-0.034	-0.385	-0.006	-0.061	0.025	0.281	0.056	0.650	0.113	1.377	0.206	2.576
R-squared	0.218		0.353		0.439		0.522		0.548		0.555		0.537		0.385	
RND volatility date t	0.654	6.894	0.716	6.567	0.693	6.247	0.605	5.205	0.552	4.496	0.503	3.933	0.404	2.920	0.125	0.698
GARCH to expiration	0.701	8.747	0.760	7.603	0.826	7.839	0.955	6.632	0.954	6.006	0.934	5.243	0.915	4.522	0.818	3.417
Historical volatility	-0.239	-4.109	-0.259	-4.121	-0.256	-3.818	-0.262	-3.552	-0.231	-3.172	-0.195	-2.557	-0.133	-1.786	-0.014	-0.155
R-squared	0.244		0.394		0.495		0.606		0.638		0.646		0.626		0.462	
GARCH to expiration	0.941	27.175	1.025	20.735	1.077	19.482	1.129	17.044	1.122	15.741	1.105	14.220	1.081	12.618	0.901	9.332
R-squared	0.229		0.370		0.469		0.583		0.617		0.628		0.615		0.460	
Historical volatility	0.583	7.457	0.636	7.842	0.667	7.872	0.690	7.630	0.689	7.858	0.683	8.010	0.676	8.149	0.570	8.126
R-squared	0.166		0.268		0.339		0.410		0.439		0.453		0.453		0.347	
GARCH to expiration	0.961	15.264	1.044	15.112	1.102	14.607	1.196	11.380	1.174	10.178	1.134	8.577	1.076	7.129	0.868	4.665
Historical volatility	-0.017	-0.299	-0.016	-0.270	-0.021	-0.333	-0.057	-0.765	-0.044	-0.576	-0.025	-0.304	0.004	0.049	0.028	0.264
R-squared	0.229		0.370		0.469		0.584		0.618		0.628		0.615		0.461	
Selected Coefficients from All-Variables Regressions																
RND volatility date t	0.567	3.585	0.607	3.758	0.599	3.779	0.482	2.520	0.382	1.775	0.293	1.350	0.125	0.514	-0.494	-1.407
GARCH to expiration	0.433	3.775	0.525	3.675	0.610	4.186	0.805	3.837	0.841	3.649	0.830	3.405	0.876	3.167	0.963	3.422
Historical volatility	-0.251	-3.915	-0.276	-3.912	-0.260	-3.625	-0.240	-3.269	-0.215	-2.813	-0.175	-2.266	-0.133	-1.756	-0.149	-1.096
R-squared	0.269		0.429		0.533		0.637		0.667		0.679		0.662		0.571	

Notes: The table presents results from regressions of realized return volatility over different horizons starting at date t+1. The top portion of the table includes only RND volatility, GARCH, and historical volatility forecasts as explanatory variables. The last four lines report coefficient estimates for the three volatility variables from multivariate regressions that also include the full set of other explanatory variables.

Table 6. Regressions of RND Volatility on All Variables for Short and Long Maturity Contracts

	Maturity < 75 days		Maturity ≥ 75 days	
	coef	t-stat	coef	t-stat
Date t return variables				
Constant	10.859	4.963	8.817	3.609
Date t return	-0.229	-8.579	-0.149	-8.896
GARCH to expiration	0.331	8.087	0.396	6.339
Date t range minus absolute return	0.387	4.709	0.317	5.748
Date t left 2% tail return	0.405	3.009	0.510	4.079
Date t right 2% tail return	0.259	3.214	0.147	2.165
Historical returns variables				
Last week return	-0.005	-11.620	-0.004	-9.969
Last week absolute return	0.003	4.208	0.002	2.756
Historical volatility (65 days)	0.090	2.370	0.108	2.946
Avg daily range (past pd = option life)	2.086	5.117	2.399	3.001
Range over last 25 days	0.161	3.671	0.125	2.296
# left 2% tail events last 2 years	0.798	8.701	0.817	7.131
Avg left 2% tail return last 2 years	-0.004	-0.017	0.101	0.411
# right 2% tail events last 2 years	0.138	2.149	0.080	0.965
Avg right 2% tail return last 2 years	0.229	1.389	0.388	2.225
Risk neutralization variables				
Michigan consumer sentiment	-0.072	-4.121	-0.063	-3.326
Baker-Wurgler investor sentiment	-0.243	-0.975	-0.425	-1.472
S&P 500 P/E ratio	-0.155	-1.409	-0.042	-0.310
BAA-AAA bond yield spread	-0.176	-0.248	-0.567	-0.628
GARCH error (past pd = option life)	-0.001	-1.076	-0.003	-1.599
GARCH RMSE (past pd = option life)	-0.001	-2.076	-0.0005	-0.951
R-squared	0.935		0.945	
NOBS	4769		5383	

Notes: The table reports results from the same RND volatility regressions shown in Table 4, broken down according to option maturity.

Table 7. Regressions of RND Volatility on All Variables Over Three Subperiods

	Jan 1996 - June 2003				July 2003 - Dec 2007				Jan 2008 - Apr 2011			
	RND volatility		Realized volatility		RND volatility		Realized volatility		RND volatility		Realized volatility	
	coef	t-stat	coef	t-stat	coef	t-stat	coef	t-stat	coef	t-stat	coef	t-stat
Date t return variables												
Constant	28.132	3.670	30.592	2.132	3.381	1.032	39.436	5.215	8.049	1.928	64.445	4.641
Date t return	-0.206	-7.057	-0.119	-1.786	-0.169	-5.418	-0.023	-0.593	-0.165	-6.678	-0.056	-0.839
GARCH to expiration	0.397	7.167	0.617	4.369	0.538	9.015	0.206	1.263	0.355	8.869	0.779	2.879
Date t range minus absolute return	0.522	6.001	0.570	3.764	0.319	3.541	0.416	2.457	0.370	4.896	0.176	0.898
Date t left 2% tail return	0.381	2.557	0.098	0.290	0.221	1.970	0.540	1.868	0.243	2.032	0.018	0.034
Date t right 2% tail return	0.315	3.418	0.154	0.778	0.096	0.633	0.173	0.633	0.161	2.165	-0.107	-0.324
Historical returns variables												
Last week return	-0.0039	-8.506	-0.0034	-2.659	-0.0056	-9.833	-0.0028	-2.597	-0.0055	-9.438	0.002	0.842
Last week absolute return	0.0005	0.621	-0.0015	-1.064	0.0012	1.599	0.0006	0.452	0.0045	6.125	-0.005	-1.253
Historical volatility (65 days)	0.066	1.058	-0.141	-1.115	0.251	4.001	0.205	0.974	0.104	3.833	-0.705	-2.222
Avg daily range (past pd = option life)	3.001	4.029	4.184	2.494	3.313	2.957	-1.970	-0.605	1.035	2.162	-2.054	-0.942
Range over last 25 days	0.044	0.999	0.210	1.986	0.102	1.603	-0.186	-1.234	0.177	3.102	-0.713	-1.667
# left 2% tail events last 2 years	0.838	5.087	-1.646	-4.303	0.452	1.959	1.599	2.227	-0.375	-1.190	5.654	2.178
Avg left 2% tail return last 2 years	-1.134	-2.284	6.022	3.915	0.366	1.848	0.896	1.664	1.502	3.955	-7.441	-2.416
# right 2% tail events last 2 years	0.289	2.022	0.575	1.654	0.071	0.869	0.105	0.415	0.226	0.789	-4.056	-5.340
Avg right 2% tail return last 2 years	-1.645	-2.146	-2.266	-1.283	1.356	3.732	-0.738	-0.900	1.601	6.278	-7.731	-2.978
Risk neutralization variables												
Michigan consumer sentiment	-0.197	-4.263	0.360	4.459	-0.057	-4.123	-0.131	-3.643	0.022	0.637	-0.413	-2.689
Baker-Wurgler investor sentiment	-0.973	-2.975	-0.174	-0.185	2.122	4.108	8.772	4.308	-0.629	-0.704	1.695	0.308
S&P 500 P/E ratio	-0.652	-0.998	-7.382	-4.405	-0.146	-0.435	-3.057	-2.316	-0.312	-1.953	-0.977	-0.975
BAA-AAA bond yield spread	-3.863	-3.369	-4.051	-1.111	0.494	0.309	-2.987	-0.852	0.770	1.109	12.099	1.985
GARCH error (past pd = option life)	-0.0016	-1.165	-0.0067	-1.729	-0.0014	-0.318	-0.0037	-0.342	-0.0007	-1.096	0.0038	1.573
GARCH RMSE (past pd = option life)	-0.0014	-2.684	-0.0006	-0.418	-0.0034	-1.936	0.0016	0.262	0.0000	0.052	0.0006	0.416
R-squared	0.794		0.437		0.892		0.684		0.957		0.680	
NOBS	3523		3523		2812		2812		3811		3811	

Notes: The table reports results from the same RND volatility regressions shown in Table 4, broken down into subsamples by time period.