The Distribution of Uncertainty: Evidence from the VIX Options Market

Clemens Völkert[‡]

This version: September 9, 2012

Abstract

This paper investigates the informational content implied in the risk-neutral distribution of the VIX, a leading barometer of economic uncertainty. We extract the risk-neutral distribution from VIX option prices over the sample period from 2006 to 2011 using a non-parametric approach. We analyze the time-series behavior of the option-implied moments and assess whether the information implied in the risk-neutral distribution has predictive power. The risk-neutral distribution considerably changed shape during the financial crisis. Furthermore, risk-neutral moments contain useful information with respect to the likelihood of upward jumps in volatility. Consistent with investors disliking high levels of economic uncertainty, we find that the overall shape of the estimated volatility pricing kernel is increasing. For certain periods, there is a puzzling U-shape. The behavior of the volatility pricing kernel over time reveals that the financial crisis has affected investors' attitudes towards risk.

Keywords: Pricing kernel, risk-neutral distribution, VIX options

JEL: G12, G13

[‡] Finance Center Münster, Westfälische Wilhelms-Universität Münster, Universitätsstr. 14-16, 48143 Münster, Germany. E-mail: clemens.voelkert@wiwi.uni-muenster.de.

1 Introduction

Derivative prices contain information about market participants' views of future states of the world. The option-implied information is useful in order to understand current market conditions and potentially to forecast future market developments. In conjunction with the objective distribution, the option-implied distribution offers insights into investors' attitudes towards risk. In this paper, we extract the distribution of volatility using VIX options. We analyze the behavior of the risk-neutral distribution over time, evaluate its predictive power, and investigate the stability and monotonicity of the volatility pricing kernel.

Starting in 1993, the Chicago Board Options Exchange (CBOE) published its VIX index. It expresses the market expectations of the 30-day volatility implied in equity index options. Since 2003, the CBOE calculates the VIX with a model-free calculation method that uses the information over the entire strike price range of S&P 500 index (SPX) options.[1] Before exchange traded volatility derivatives were introduced, over-the-counter variance and volatility swaps were the most direct way to trade volatility. VIX futures were listed at the CBOE Futures Exchange (CFE) in 2004, VIX options started trading at the CBOE two years later.

Options on the VIX are particularly suited to study market disruptions. When stock prices decline the VIX usually increases dramatically. Due to this inverse relationship between stock prices and volatility, investors consider the VIX as a leading barometer of uncertainty in the market and commonly call it "investor fear gauge".[2] Bloom (2009) points out that stock market volatility is linked to productivity and demand uncertainty. In this sense, using VIX options it is possible to extract the distribution of economic uncertainty. The volume and open interest of VIX derivatives

[1]The "old" VIX (now VXO) was based on S&P 100 index at-the-money put and call prices and calculated using the Black and Scholes (1973) option pricing model. For a detailed description of the "new" VIX see Jiang and Tian (2007), CBOE (2009), and Whaley (2009).

[2]See Whaley (2000) and Low (2004). Other option-implied measures of "fear" are discussed in Bollerslev and Todorov (2011) and Schneider (2012).

have increased rapidly since their introduction. By now, VIX options are the second most actively traded contract at the CBOE. The high liquidity ensures reliable options data to construct risk-neutral distributions.

The papers closest to ours from a methodological point of view are Bliss and Panigirtzoglou (2002, 2004) and Kostakis et al. (2011). We construct constant maturity risk-neutral densities using a non-parametric method. A smooth implied volatility surface is obtained by first interpolating implied volatilities across options with a given time to expiration and then in the maturity dimension. Afterwards, the result in Breeden and Litzenberger (1978) is used to extract the risk-neutral density. The papers above apply this method to equity index options, while we focus on volatility options. To the best of our knowledge, this is the first study which extracts the risk-neutral distribution from VIX option prices and analyzes its behavior over time.

There is a large body of literature about the informational content of option-implied densities from equity indices, interest rates, and exchange rates. Many of these studies focus on the second moment of the option-implied distribution. Canina and Figlewski (1993), Christensen and Prabhala (1998), and Jiang and Tian (2005) use the implied volatility to forecast the future realized volatility of the underlying asset. There is strong empirical evidence that implied volatilities have predictive power and are superior to time-series based volatility forecasts. Other studies focus on the shape of the risk-neutral distribution around specific events, like elections and market crashes, and assess whether these events are predictable or not. Examples include Melick and Thomas (1997) and Coutant et al. (2001). Bates (1991) investigates S&P 500 futures options over the period from 1985 to 1987 and shows that the 1987 stock market crash was anticipated by the options market. Subsequent research has found inconclusive evidence concerning the predictability of market crashes using risk-neutral distributions.[3]

[3] A review of the literature is given by Christoffersen et al. (2012).

There are several studies which analyze the financial crisis based on stock market volatility and equity index option prices. Schwert (2011) investigates stock market volatility during the recent crisis and compares it with previous periods of high volatility. He points out that compared to the Great Depression the recent crisis was relatively short-lived. Birru and Figlewski (2012) analyze the events surrounding the collapse of Lehman Brothers using real time SPX options data. Remarkably, the risk-neutral skewness and kurtosis did decrease in magnitude. This indicates that drops in the S&P 500 index were not perceived to be more likely.[4] In contrast to these studies, we analyze the risk-neutral volatility distribution. We find that the option-implied volatility distribution considerably changed shape over the sample period. Visual inspection of the risk-neutral distribution shows that it primarily reacts to major market disruptions like the collapse of Lehman Brothers. We employ a binary response model to forecast upward spikes in volatility and find that for short times to maturity option-implied moments contain useful information with respect to the likelihood of upward jumps in the VIX.

Option prices incorporate preferences towards risk as well as beliefs about future outcomes. The discounted ratio between the risk-neutral distribution and its objective counterpart is the pricing kernel. Using equity index options data, Jackwerth (2000), Rosenberg and Engle (2002), and Bakshi et al. (2010) show that in contrast to standard economic theory, the pricing kernel is not monotonically declining in market returns. In subsequent research, the puzzling behavior of projected pricing kernels has been explored for interest rates and certain commodities.[5] Economic intuition suggests that investors dislike economic uncertainty and thus, the pricing kernel should be increasing in the volatility dimension. Bakshi et al. (2012) find evidence for a non-monotonicity in the volatility pricing kernel and propose a model with heterogeneous agents that features increasing and decreasing regions of

[4] Neumann and Skiadopoulos (2012) find that the implied skewness and kurtosis did not change significantly during the financial crisis.

[5] See e.g. Beber and Brandt (2006), Li and Zhao (2009), and Pan (2012).

the volatility pricing kernel. We find that the overall shape of the estimated volatility pricing kernel is increasing, though some of the estimates are locally decreasing. Furthermore, there is considerable variation in the shape of the volatility pricing kernel over the sample period. This indicates that the financial crisis has affected investors' attitudes towards risk.

The remainder of this paper is organized as follows. In Section 2, we introduce the methodology to construct constant maturity risk-neutral densities and estimate the volatility pricing kernel. In Section 3, we analyze the option-implied information, assess the probability the market attaches to upward spikes in future values of the VIX, and investigate the shape of the volatility pricing kernel. Section 4 concludes.

2 Empirical Methodology

2.1 Derivatives on the CBOE volatility index

In September 2003, the CBOE revised its volatility index. The new calculation methodology makes it feasible to replicate the VIX based on a portfolio of SPX options. A major reason for the revision of the VIX was to create a suitable underlying for tradable volatility products. VIX futures and options are contracts on the forward 30-day implied volatility of SPX options. VIX futures started trading on March 26, 2004. European style options on the VIX followed on February 24, 2006. VIX options expire on the Wednesday that is thirty days prior to the third Friday of the calendar month immediately following the expiring month. The CBOE lists up to six contract months, provided that the time to expiration is no greater than 12 months. VIX options are cash-settled based on a special opening quotation calculated from a sequence of opening prices of the options used to calculate the VIX on the settlement date.

VIX futures and options data are obtained from the CBOE's website and Market Data Express, respectively. The sample period ranges from February 24, 2006 to December 31, 2011. We apply several filters to the options data. We eliminate option quotes that do not satisfy standard no-arbitrage conditions. We also check for negative bid-ask spreads, zero bids, and filter out all options with zero open interest and those options where the implied volatility could not be computed. We exclude in-the-money (ITM) options (puts with moneyness greater than 1.05 and calls with moneyness smaller than 0.95) because they are less liquid than out-of-the-money (OTM) and at-the-money (ATM) options. Moneyness is defined as the strike price divided by the VIX futures price. Options with time to maturity shorter than 7 days are discarded to reduce pricing anomalies that might occur close to expiration. Contracts with maturity of more than one year are also excluded.[6] We use the mid of the bid and ask prices in the following. Table 1 displays the options data separated into moneyness-maturity categories. For VIX futures, we employ daily settlement prices. Constant maturity Treasury bill yields are used for the risk-free discount rate. The risk-free rate for a particular time to maturity is obtained by linearly interpolating the two discount rates which straddle the maturity of the options.

2.2 Extracting the risk-neutral distribution

There are several methods that can be employed to extract risk-neutral distributions from option prices. Jackwerth (1999, 2004) provides a literature review and separates the approaches into two basic categories. Parametric methods use a probability distribution, for instance a mixture of lognormals as in Melick and Thomas (1997), and calibrate the parameters by matching observed option prices. Non-parametric approaches first fit a flexible function of option prices or implied volatilities across strike

[6]Currently, there are no VIX options with maturity beyond one year. Prior to the last changes in the contract specifications, up to three near-term months and up to three additional months on the February quarterly cycle were listed.

prices and then use the method of Breeden and Litzenberger (1978) to extract the risk-neutral distribution. Bliss and Panigirtzoglou (2002) find that non-parametric methods are easier to handle computationally and also more stable than commonly used parametric methods. Because of these advantages, we use a non-parametric method and fit a function of implied volatilities across deltas.

Breeden and Litzenberger (1978) show that the second derivative of a European call option price with respect to its strike price is the discounted risk-neutral probability of the future asset price ending up at exactly the strike price of the option. The price of a VIX call option with strike price K and maturity in T is

$$C_t(K,T) = \int_K^\infty e^{-\int_t^T r_s ds} \left(\mathrm{VIX}_T - K\right) q_t(\mathrm{VIX}_T, T)\, d\mathrm{VIX}_T,$$

where q denotes the risk-neutral density and r is the risk-free rate. Differentiating twice with respect to the strike price and rearranging yields

$$q_t(\mathrm{VIX}_T, T) = e^{\int_t^T r_s ds} \left. \frac{\partial^2 C_t(K,T)}{\partial K^2} \right|_{K=\mathrm{VIX}_T}. \tag{1}$$

Equation (1) requires a continuum of traded option prices. However, in practice we typically observe relatively few option prices for discretely spaced strike prices. Furthermore, options on the VIX have fixed expiry dates. Thus, when time passes the time to maturity decreases and this makes it difficult to compare option-implied distributions over time. We eliminate the discreteness and the time to maturity biases and construct constant maturity densities following Bliss and Panigirtzoglou (2002, 2004) and Kostakis et al. (2011). Essentially, a smooth implied volatility surface is constructed by first interpolating across options with a given time to expiration and then in the maturity dimension. Afterwards Equation (1) is used to extract the risk-neutral distribution using finite difference methods.

Implied volatilities are calculated using the Black (1976) formula, i.e. the implied volatility of VIX options is the volatility parameter that, plugged into the Black (1976) option pricing formula, makes market and model prices agree. To reduce the effect of a jump at the transition point between call and put options, we blend the call and put implied volatilities in the region around the ATM level in a similar way as Figlewski (2010).[7] Thus, we use put option implied volatilities for moneyness below 0.95, blended implied volatilities for the interval [0.95 1.05], and call option implied volatilities for moneyness above 1.05.

Following Malz (1997) and Bliss and Panigirtzoglou (2002, 2004), the inter- and extrapolation is performed in the delta rather than the strike price space. A major advantage of this choice is that the delta of a call option is bounded between zero and one. To convert implied volatilities from the strike price space to the delta space, we use the Black (1976) formula with the ATM implied volatility. Using the ATM implied volatility ensures that the ordering of the deltas is the same as the ordering of the strike prices. We interpolate implied volatilities across delta using a cubic smoothing spline and extrapolate beyond the maximum (minimum) available strike price with the implied volatility of the highest (lowest) actually traded strike price. The smoothing parameter is initialized at 0.99 and, if necessary, adjusted upward to fit option prices within their bid-ask spreads. We use a constant delta span from 0.01 to 0.99 for all contracts. The fitted spline leaves us with a narrowly spaced set of implied volatilities across deltas.

For each time to maturity, we extract 99 deltas ranging from 0.01 to 0.99. We interpolate in the time dimension to construct a hypothetical implied volatility smile for fixed times to maturity (30, 60, and 120 days). We fit a cubic spline through the constant maturity implied volatilities and extract 20,000 data points for deltas

[7] Whenever there are quotes for call and put options with the same moneyness M in the interval [0.95 1.05], a blended value of the implied volatility is computed as

$$\mathrm{IV}_{blend}(M) = w\ \mathrm{IV}_{put}(M) + (1-w)\ \mathrm{IV}_{call}(M), \quad \text{where} \quad w = \frac{1.05 - M}{1.05 - 0.95}.$$

between 0.01 and 0.99. In the next step, the implied volatilities are converted into call prices using the Black (1976) option pricing model. This fine grid of prices is used to approximate the risk-neutral distribution in Equation (1) with finite difference methods. We often express the density in terms of log returns $R_t = \log\left(\frac{VIX_T}{F_t(T)}\right)$, where $F_t(T)$ denotes the VIX futures price at time t with maturity in T. By a change of variable, it follows that $q_t(R,T) = \text{VIX}_T\, q_t(\text{VIX}_T, T)$. The estimated risk-neutral densities are normalized so that they integrate to one.

To compute summary statistics, we numerically integrate the appropriate function of the probability density to estimate the moments. More specifically, the k-th central moment can be calculated using

$$m_k(x) = \int_{-\infty}^{\infty} (x-\mu)^k\, q(x)\, dx, \quad \text{with} \quad \mu = \int_{-\infty}^{\infty} x\, q(x)\, dx.$$

The variance can be obtained by setting k equal to 2. Skewness and kurtosis follow from $\frac{m_3}{m_2^{3/2}}$ and $\frac{m_4}{m_2^2}$, respectively. We use the cumulative distribution function to estimate selected percentiles.

2.3 Estimating the volatility pricing kernel

The "true" pricing kernel depends on multiple economic factors, for example aggregate wealth, interest rates, volatility, etc. By using options, it is possible to compute the projected pricing kernel, i.e. the pricing kernel projected onto the payoffs of the underlying asset. Jackwerth (2000) and Rosenberg and Engle (2002) project the pricing kernel onto equity returns, while Beber and Brandt (2006) and Li and Zhao (2009) investigate the pricing kernel obtained from interest rate derivatives. Each market provides a subset of information about the "true" pricing kernel and, depending on the underlying asset, it is possible to study preferences from different perspectives.

Using VIX options data, we obtain an estimate of the volatility pricing kernel. It is the discounted ratio between the risk-neutral and the objective VIX densities

$$\xi_t = e^{-\int_t^T r_s ds} \, \frac{q_t(\mathrm{VIX}_T, T)}{p_t(\mathrm{VIX}_T, T)}.$$

Similar to Rosenberg and Engle (2002) and Barone-Adesi et al. (2008), we estimate objective probabilities using an ARMA-GARCH time-series model. Instead of modeling the VIX directly, we employ the logarithmic transformation. Andersen et al. (2001) show that the distribution of log volatility is closer to normal. More specifically, we select a suitable time-series model following the Box-Jenkins methodology. We find that an ARMA(2,1)-GARCH(1,1) model fits the data very well. For each day, we estimate the time-series model based on the past 10 years of VIX data. Afterwards, we simulate the time-series model 100,000 times and use a kernel density estimator to obtain an estimate of the objective probability density function. We employ a normal kernel with standard bandwidth, i.e. $0.9N^{-1/5}\sigma$, where N is the number of simulations and σ is the standard deviation of the simulated values. The probability distributions are scaled to make them integrate to one.

3 Results

3.1 The VIX over the sample period

We start with a description of the major events during and around the recent financial crisis.[8] Figure 1 displays the VIX and the realized volatility from February 24, 2006 to December 31, 2011.[9] Table 2 reports the largest 35 daily percentage increases

[8]A timeline of the financial crisis is available at the St. Louis Fed webpage (`http://timeline.stlouisfed.org/index.cfm?p=timeline&tlid=538#tl_538`).

[9]Realized volatilities are calculated by summing up squared 5-minute intra-day returns for each day and afterwards taking the square-root. Our data source is the Oxford-Man Institute's "realized library" (`http://realized.oxford-man.ox.ac.uk/`).

in the VIX. It also shows the returns to the S&P 500 index. Large increases in the VIX are accompanied by sharp drops in the S&P 500 index. The financial crisis was associated with unusually high levels of volatility, while the period preceding it displayed low levels of volatility. In 2005 and 2006 the average closing values of the VIX were below 13, while they reached more than 30 in 2008 and 2009. To put this into perspective, the mean level of the VIX from 1990 to 2011 was 20.

The first sign of the upcoming period of high volatility was signaled by the "Shanghai surprise" on February 27, 2007 (VIX +64.22%, SPX -3.47%). Close to 10% drops in major Chinese stock market indices triggered the biggest one day percentage increase in the VIX ever recorded. Less than half a year later on August 9, 2007 BNP Paribas announced that it halts redemptions on three investment funds (VIX +23.45%, SPX -2.96%). According to the National Bureau of Economic Research (NBER), the U.S. economy went into a recession in December 2007. However, the financial crisis did not immediately impact the stock market. A strong decrease in leading U.S. stock price indices happened in the fall of 2008, accompanied by a massive increase in economic uncertainty measured by the VIX. Lehman Brothers declared bankruptcy on September 15, 2008 (VIX +23.54%, SPX -4.71%). On the same day, Merrill Lynch was sold to Bank of America. On September 29, 2008 the U.S. House of Representatives voted down the Troubled Asset Relief Program (TARP) bailout plan leading to high levels of uncertainty in the market (VIX +34.48%, SPX -8.79%). Shortly after its implementation (Emergency Economic Stabilization Act on October 3, 2008), several large financial institutions announced their intentions to subscribe to the program due to the diminished liquidity in the interbank funding market. The VIX reached its historical high of 80.86 on November 20, 2008. With a close to zero federal funds target rate, tighter money market spreads, and strong fiscal stimulus, volatility resumed to moderate levels in 2009.

The decline in economic activity and the need to finance the fiscal stimulus led to rising levels of government debt across the globe and shifted the concerns from the private to the public sector. Rating downgrades in some European countries generated renewed concerns about the stability of the financial system. The sovereign debt crisis commenced in late 2009. The "flash crash" on May 6, 2010 (VIX +31.67%, SPX -3.24%), where the VIX reached its highest level since April 2009, can be regarded as a sign of market instability. In 2011, the high levels of debt in the U.S. increased the anxiety of a sovereign default due to breaching the debt ceiling. With several of the largest percentage increases in the VIX, August 2011 was the most volatile month since the fall of 2008. Standard & Poors downgraded the U.S. government credit rating on August 5, 2011 from AAA to AA+ due to large deficits and increasing levels of debt. From August 3, 2011 to August 8, 2011 the VIX more than doubled. Following the downgrade, weak U.S. economic indicators and European banking concerns on August 18, 2011 amplified uncertainty (VIX +35.12%, SPX -4.46%). By the end of 2011, markets recovered with the VIX closing slightly above 20, its long-run average.

3.2 Time-series behavior of option-implied moments

To focus on the major properties of the risk-neutral distribution, we extract 30-, 60-, and 120-day risk-neutral moments. First, we investigate the risk-neutral mean and the 90% confidence interval. Afterwards, we take a closer look at the behavior of the higher moments.

The mean of the risk-neutral distribution and the 90% confidence interval are displayed in Figure 2. The risk-neutral mean coincides with the VIX futures price. Over the sample period it was on average about 24. The risk-neutral mean is less volatile than the spot VIX and the variability decreases with increasing time to maturity. One measure of uncertainty is the width of the confidence interval. The dotted vertical lines in Figure 2 correspond to the following events (in chronological order): 1. Shanghai surprise, 2. BNP Paribas stops redemptions on some of its funds, 3. Lehman Brothers files for Chapter 11 bankruptcy, 4. TARP failure, 5. 2010 flash crash, and 6. Downgrade of U.S. credit rating. All but one of these events, the Shanghai surprise, were succeeded by periods of high uncertainty in the market.

The Shanghai surprise represents the largest daily percentage increase in the VIX over the sample period. However, the width of the confidence band did not increase significantly. This is plausible, since there were hardly any economic or political reasons for the large rise in volatility. In contrast, the BNP Paribas incident resulted in a considerably wider gap between the 95th and 5th percentiles, which persisted until September 2008. The most notable increase in the width of the confidence band occurred during the market turbulence in the fall of 2008. It was likely due to the increased uncertainty what policy makers would do after Lehman Brothers filed for Chapter 11 bankruptcy protection and the rejection of TARP by the House of Representatives. Even during this turbulent period, the VIX options market did not indicate volatility levels above 100 (95th percentile). On November 20, 2008, the date where the VIX reached its historical high of 80.86, the risk-neutral mean values were considerably lower (65 for 30 days, 60 for 60 days, and 51 for 120 days). This shows that the extremely high levels of implied volatility were not expected to last for long periods of time.

The uncertainty, measured by the width of the confidence band, remained large for several months. Due to heavy fiscal and monetary interventions, markets recovered in 2009, volatility gradually reverted back to its mean level, and the confidence interval tightened to some extent. Just before the 2010 flash crash, the 95th and 5th percentiles were back at their mid 2008 values and the VIX was down at its long-term average. However, markets were extremely fragile due to the sovereign debt crisis, in particular the instability of Greece. The flash crash resulted in a jump in the risk-neutral mean of the VIX and an extreme increase in the width of the confidence band. Mean-reversion pulled the level of volatility back to about 20 by the end of 2010. The downgrade of the U.S. government credit rating on August 5, 2011 triggered another sharp increase in the width of the confidence band, which persisted until the end of the sample period. Overall, we find that the risk-neutral mean and the width of the confidence interval suggest that several key events triggered the high levels of uncertainty in the market.

Table 3 displays summary statistics of the risk-neutral moments in levels (Panel A) and first differences (Panel B). Figures 2 and 3 show the time-series of the option-implied moments. Similar to the risk-neutral mean, the risk-neutral variance went up significantly during the turbulent times in the recent financial crisis and then mean-reverted to its pre-crisis level. Furthermore, we observe a pronounced right skewness and excess kurtosis. The right skewness indicates that market participants attach a high probability to upward movements in volatility. The excess kurtosis shows that the tails of the risk-neutral volatility distribution are relatively heavy.

Table 4 shows correlations between the risk-neutral moments. The correlation between the risk-neutral mean and variance is significantly positive, ranging from 0.83 to 0.88 with increasing time to maturity. Furthermore, the risk-neutral skewness and kurtosis have a very strong positive linear association. The cross-correlations between the other moments are more moderate and decreasing in the time horizon. Table 5 displays autocorrelations of the risk-neutral moments in levels (Panel A) and first differences (Panel B). The decay in the autocorrelations is very slow and we observe high values even for 50 lags. This casts doubt in the stationarity of the time-series. To test if the time-series are stationary, we employ the Phillips-Perron unit root test. Except for the risk-neutral mean, the test rejects non-stationarity. Thus, even though the higher risk-neutral moments are persistent, they gradually revert to their long-run averages. The autocorrelations are higher for longer time horizons. All series are stationary in their first differences. Note that the large negative first order autocorrelations of the first-differenced risk-neutral skewness and kurtosis indicate that the higher moments are predictable. It might be possible to use option trading strategies to capitalize on the predictability of the higher risk-neutral moments.[10]

3.3 Relation between VIX and risk-neutral distribution

There is merely a shift in the risk-neutral distribution if an increase in the VIX is associated with an identical increase in each quantile of the risk-neutral distribution. To assess the connection between the VIX and the shape of the distribution, we follow Figlewski (2010) and regress the quantiles of the risk-neutral distribution onto the VIX. Due to the high autocorrelation in the levels, we use first differences.

The results are summarized in Table 6. The slope coefficients are all significant and they increase in the confidence level, meaning that a change in the VIX impacts the right tail of the option-implied distribution more severely than the center or the left tail. The low sensitivity of the left tail is intuitive, since the VIX is bounded from bellow. For longer horizons, the shape of the distribution responds more moderately to changes in the VIX. The typical long right tail of the risk-neutral distribution can be attributed to the possibility of large upward jumps in the VIX. When the VIX reaches high levels, mean-reversion pulls it back. Thus, after a significant increase in volatility, we expect (ceteris paribus) skewness to gradually decrease over time. We anticipate the same behavior for the risk-neutral kurtosis. As the risk-neutral moments are very persistent, this effect requires some time to unfold.

[10]The predictability of higher order risk-neutral moments inferred from SPX options has been investigated by Neumann and Skiadopoulos (2012). They find that it is not possible to earn abnormal profits, after controlling for transaction costs.

3.4 Predictability

3.4.1 Risk-neutral distribution around major events

We compare the estimated risk-neutral densities before and after the six events outlined in Section 3.2. Figure 4 and Table 7 indicate that market participants revise their expectations concerning future values of volatility after major incidents. Table 7 shows the average values of the VIX and the risk-neutral moments 30 days before and after each event. The most notable finding is that the VIX, the risk-neutral mean, and the risk-neutral variance increase dramatically after each incident. There is no clear pattern for the risk-neutral skewness and kurtosis. We pointed out in Section 3.3 that the risk-neutral skewness and kurtosis are likely to decrease after significant increases in volatility due to mean-reversion. Whether the higher moments actually decrease or increase is a trade-off between volatility moving to its average value over time and the impact of news. After the Shanghai surprise and the downgrade of the U.S. credit rating, the 30-day risk-neutral skewness decreased. Subsequent to the other events, the 30-day risk-neutral distribution was more right-skewed, which reflects a high probability attached to large future values of the VIX.

Figure 4 displays risk-neutral distributions one day before and after each event. The densities have a fixed time horizon of 30 days. The risk-neutral distributions for longer horizons have similar shapes. After each event, the risk-neutral distribution is typically wider with a longer right tail. This implies that market participants expected further increases in the VIX. The smallest differences in the shape of the risk-neutral density occurred for the Lehman crash. Compared to the other events, the Lehman crash was associated with a relatively small percentage increase in the VIX. Figure 4 shows that views of future implied volatility did not change directly after the collapse of Lehman Brothers.[11] This reveals that market participants needed

[11]Birru and Figlewski (2012) analyze the risk-neutral S&P 500 index distribution around the Lehman crash and find that the higher moments did not change as anticipated.

some time to revise their views of future market volatility after the crash. By the end of September 2008, investors were expecting significantly higher volatilities.

3.4.2 Forecasting jumps in volatility

We employ a probit model to forecast upward spikes in the VIX

$$\text{Prob}(Y = 1|x) = \Phi(x'\beta),$$

where Φ denotes the standard normal distribution and x is a vector of explanatory variables. The coefficients β reflect the impact of changes in the explanatory variables on the probability. The probit model has been extensively used to forecast recessions and market crashes.[12] Doran et al. (2007) find that the risk-neutral skewness of the S&P 100 index has predictive power in forecasting market declines. While they investigate large changes in stock returns, we analyze the predictability of upward spikes in volatility. The binary variable Y is 1 if there is a large increase in the VIX over the life of the option and zero otherwise. We define a volatility spike as either the 97.5% or the 99% quantile of the empirical distribution of VIX log returns.

Due to the high cross-correlations between the risk-neutral moments, we only use the risk-neutral variance and skewness as explanatory variables. A priori, we expect both coefficients to be positive, i.e. we anticipate a higher implied variance and a more pronounced right skewness prior to upward spikes in volatility.[13] We control for the effect of trading pressure as in Dennis and Mayhew (2002) and Beber and Brandt (2006) using put to call ratios based on open interest and volume. Trading volumes and open interest of VIX options increased rapidly after their introduction. As market participants use OTM VIX call options to protect their portfolios

[12] See e.g. Estrella and Mishkin (1998), Kumar et al. (2003), and Doran et al. (2007).

[13] We investigated other common measures of skewness, e.g. the spread between OTM and ATM implied volatilities. The results are similar to the ones reported below.

against sharp decreases in stock prices and increases in volatility, call options are more heavily traded compared to put options. If the put to call ratio decreases, more calls relative to puts are bought. This constitutes a hedge against increasing volatility. Consequently, the expected signs of the put to call ratio coefficients are negative. We also include the yield spread, a prominent predictor of real economic activity, see for instance Estrella and Mishkin (1998). It represents the difference between the yields of 10-year and 3-month Treasury securities. The yield data is obtained from the St. Louis Fed. A narrow yield spread or an inverted yield curve is usually associated with poor economic conditions. As volatility is high when economic conditions are bad, this shape of the yield curve might signal an increased probability of an upward spike in volatility.

We focus on a 30-day horizon of the risk-neutral distribution of log returns. Longer term options have less predictive power in forecasting upward spikes in volatility. The model is estimated using maximum likelihood. We account for the autocorrelation due to overlapping observations using Newey and West (1987) standard errors. The results are displayed in Table 8.

The average values of the explanatory variables if an upward jump in the VIX occurred during the life of the option for a 97.5% (99%) cutoff level are: 0.0495 (0.0516) for risk-neutral variance and 0.6551 (0.7071) for risk-neutral skewness. If the dependent variable takes a value of zero, the average values are lower: 0.0407 (0.0421) for risk-neutral variance and 0.6366 (0.6319) for risk-neutral skewness. The put to call ratios indicate that calls are more actively traded than puts. However, the ratios are not lower prior to upward spikes in volatility. The mean yield spread is larger prior to jumps in volatility: 2.2588 (2.3006) versus 1.9364 (1.9944).

The estimated coefficients of the risk-neutral variance and skewness are 18.9788 (17.0284) and 0.7859 (1.3103). Both coefficients have the expected sign. An increase in either of the variables leads to a rise in the predicted probability. While the risk-neutral variance is significant for both cutoff levels, the risk-neutral skewness shows no significant relation to the probability of an upward spike in volatility for the 97.5% volatility jump definition. The coefficients of the put to call ratios are not statistically different from zero. Thus, we find no evidence for the hypothesis of increased trading of call options relative to put options prior to upward spikes in volatility. The yield spread is also not related to the likelihood of large increases in volatility. The joint hypothesis that the slope coefficients are all equal to zero can be tested using the likelihood ratio test. The test statistic is: $-2\left(\ln L_0 - \ln L\right) = 86.4024$ (72.1526). The χ^2 distribution with 5 degrees of freedom and a critical value of 5% is 11.0705, so the joint hypothesis that all slope coefficients are equal to zero is rejected. Similarly, the hypothesis that only the coefficients of the option-implied moments are equal to zero can be rejected. These results indicate that the information implied in VIX options is a leading indicator of upward spikes in the VIX.

3.5 Volatility pricing kernel

Option prices incorporate preferences towards risk as well as beliefs about future outcomes. Changes in the risk-neutral distribution over time can thus either be due to revisions of market participants' attitudes towards risk or changes in the objective distribution. The ratio of the risk-neutral and the objective distribution, the pricing kernel, provides an assessment of how market participants assess risk. By tracking the behavior of the pricing kernel over time, it is possible to investigate whether investors' attitudes towards risk changed. The shape of the projected pricing kernel has been investigated for several underlyings. For stock indices, interest rates, and certain commodities there is evidence for a U-shape.[14]

We use the estimated risk-neutral distribution and the objective distribution

[14] See e.g. Jackwerth (2000), Rosenberg and Engle (2002), Beber and Brandt (2006), Li and Zhao (2009), Bakshi et al. (2010), and Pan (2012).

from an ARMA-GARCH model to obtain the pricing kernel projected onto volatility. We focus on a 30-day time horizon.[15] To visualize the results, we map the physical density and the risk-neutral density on a fine grid of log returns. Figure 5 displays the estimates of the risk-neutral density, the objective density, and the volatility pricing kernel. If investors are risk-averse, risk premia drive a wedge between risk-neutral and objective probabilities attached to future levels of volatility. We find that all moments of the objective distribution are lower. The risk-neutral density assigns more weight to periods of high volatility than its objective counterpart. Thus, the volatility pricing kernel is upward sloping in the right tail. Investors attach high values to payoffs received when aggregate wealth is low. High levels of volatility tend to coincide with poor economic conditions and an upward sloping volatility pricing kernel implies that volatility moves in the opposite direction as aggregate wealth.

Bakshi et al. (2012) test the monotonicity of the volatility pricing kernel. They construct time-series of non-overlapping option returns for different moneyness categories and investigate whether the differences between the average returns are significant or not. They find that VIX put option returns are increasing in the strike price. This indicates that the volatility pricing kernel is downward sloping in the left region. We observe the same behavior for certain periods, especially at the beginning of our sample. However, on average the center and the left tail of the volatility pricing kernel in Figure 5 are rather flat than decreasing. A non-monotonic pricing kernel reflects differential pricing of downside and upside risks. A downward sloping left tail implies that market participants regard low values of the VIX as bad states. McMillan (2004, Ch. 4) points out that extremely low levels of the VIX are not necessarily a good sign, as very low prices of equity index options might signal an upcoming stock market correction in either direction ("calm before the storm").

[15]For longer times to maturity the estimated volatility pricing kernels have similar shapes. The results are robust to a shorter estimation window for the time-series model and the empirical innovations density proposed by Barone-Adesi et al. (2008).

To investigate the behavior over time, we plot the average volatility pricing kernel for each year of the sample period in Figure 6. There is substantial variation in the shape of the volatility pricing kernel over the sample period. While the right part of the volatility pricing kernel is always upward sloping, the center and the left tail show different shapes. In 2006, 2007, and 2011 Figure 6 suggests a non-monotonic volatility pricing kernel, which is consistent with the finding in Bakshi et al. (2012). In particular, 2006 was characterized by unusually low levels of the VIX. Following McMillan's argument, the U-shape shows that market participants perceived that very low levels of the VIX precede a declining stock market.

The shift in the level and shape of the volatility pricing kernel could be due to the recession (December 2007 to June 2009) in our sample period. Among others Rosenberg and Engle (2002) and Bollerslev et al. (2011) find that risk aversion is countercyclical. Furthermore, Boyarchenko (2012) shows that the amount of ambiguity faced by investors changed during the financial crisis. If investors associate high levels of volatility with states of the world in which their marginal utility is high, assets that pay off in these scenarios are valuable. Volatility related derivatives, for instance variance swaps, pay off in these states, leading the long party to accept a negative premium. For higher levels of risk and ambiguity aversion the insurance against high levels of volatility is more valuable. Thus, it is anticipated that the volatility premium also displays a countercyclical behavior. As a crude estimate of the volatility premium, we use the difference between the risk-neutral and realized volatilities in Figure 1.[16] The median difference between the VIX and the realized volatility for each year are: -4.01 (2006), -5.00 (2007), -6.32 (2008), -10.68 (2009), -8.91 (2010), and -7.80 (2011). The large gap between the two in 2009 and 2010 is consistent with the steep right tail of the volatility pricing kernel in Figure 6. The behavior of the volatility pricing kernel over time indicates that there is a consid-

[16]The existence of a volatility premium can also be inferred from Figure 5, which shows that the risk-neutral density is shifted to the right relative to its objective counterpart.

erable shift in how market participants price volatility. They put more weight on high volatility states in the recession period. In addition, during this time period the volatility pricing kernel shows no U-shape. In 2011, the shape is similar to the one in 2006 and 2007. One explanation for this behavior is that low levels of volatility during and shortly after the recession period were regarded as unambiguously favorable signs, while this was not the case prior to the financial crisis and in 2011.

4 Conclusion

One important source of information about market participants' perception of aggregate stock market uncertainty is the volatility implied in equity index option prices. We extract the risk-neutral distribution of the VIX from 2006 to 2011 using a non-parametric method. The risk-neutral volatility distribution significantly changed shape over the sample period. Consistent with the evidence for equity index options in Gemmill and Saflekos (2000) and Lynch and Panigirtzoglou (2008), we find that the risk-neutral distribution primarily reacts to major market disruptions like the collapse of Lehman Brothers or the downgrade of the U.S. credit rating. We find evidence that the option-implied variance and skewness contain useful information with respect to the likelihood of upward jumps in the VIX.

Consistent with economic theory, we find that states of the world in which volatility is high are the ones in which investors have high marginal utility. While an upward sloping volatility pricing kernel is consistent with investors disliking high levels of uncertainty, the shape of the left tail is puzzling for certain periods of time. The behavior of the volatility pricing kernel over time shows that investors' attitudes towards risk changed considerably. The higher values of the volatility pricing kernel during and after the recession show that protection sellers commanded higher insurance premia for large levels of volatility when economic conditions were bad.

References

Andersen, T., T. Bollerslev, F. Diebold, and H. Ebens (2001): "The Distribution of Realized Stock Return Volatility," *Journal of Financial Economics*, 61, 43–76.

Bakshi, G., D. Madan, and G. Panayotov (2010): "Returns of Claims on the Upside and the Viability of U-Shaped Pricing Kernels," *Journal of Financial Economics*, 97, 130–154.

——— (2012): "Heterogeneity in Beliefs and Volatility Tail Behavior," Working Paper.

Barone-Adesi, G., R. Engle, and L. Mancini (2008): "A GARCH Option Pricing Model with Filtered Historical Simulation," *Review of Financial Studies*, 21, 1223–1258.

Bates, D. (1991): "The Crash of '87: Was It Expected? The Evidence from Options Markets," *Journal of Finance*, 46, 1009–1044.

Beber, A. and M. Brandt (2006): "The Effect of Macroeconomic News on Beliefs and Preferences: Evidence from the Options Market," *Journal of Monetary Economics*, 53, 1997–2039.

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.

Black, F. (1976): "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3, 167–179.

Black, F. and M. Scholes (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–659.

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

——— (2004): "Option-Implied Risk Aversion Estimates," *Journal of Finance*, 59, 407–446.

Bloom (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

Bollerslev, T., M. Gibson, and H. Zhou (2011): "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities," *Journal of Econometrics*, 160, 235–245.

Bollerslev, T. and V. Todorov (2011): "Tails, Fears, and Risk Premia," *Journal of Finance*, 66, 2165–2211.

Boyarchenko, N. (2012): "Ambiguity Shifts and the 2007-2008 Financial Crisis," *Journal of Monetary Economics*, 59, 493–507.

Breeden, D. and H. Litzenberger (1978): "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business*, 51, 621–651.

Canina, L. and S. Figlewski (1993): "The Informational Content of Implied Volatility," *Review of Financial Studies*, 6, 659–681.

CBOE (2009): "The CBOE Volatility Index – VIX," White Paper.

Christensen, B. and N. Prabhala (1998): "The Relation between Implied and Realized Volatility," *Journal of Financial Economics*, 50, 125–150.

Christoffersen, P., K. Jacobs, and B. Chang (2012): "Forecasting with Option Implied Information," in *Handbook of Economic Forecasting*, ed. by G. Elliott and A. Timmermann, Elsevier, vol. 2.

Coutant, S., E. Jondeau, and M. Rockinger (2001): "Reading PIBOR Futures Options Smiles: The 1997 French Snap Election," *Journal of Banking and Finance*, 25, 1957–1987.

Dennis, P. and S. Mayhew (2002): "Risk-Neutral Skewness: Evidence from Stock Options," *Journal of Financial and Quantitative Analysis*, 37, 471–493.

Doran, J., D. Peterson, and B. Tarrant (2007): "Is There Information in the Volatility Skew?" *Journal of Futures Markets*, 27, 921–959.

Estrella, A. and F. Mishkin (1998): "Predicting U.S. Recessions: Financial Variables as Leading Indicators," *Review of Economics and Statistic*, 80, 45–61.

Figlewski, S. (2010): "Estimating the Implied Risk Neutral Density for the U.S. Market Portfolio," in *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, ed. by T. Bollerslev, J. Russell, and M. Watson, Oxford: Oxford University Press.

Gemmill, G. and A. Saflekos (2000): "How Useful are Implied Distributions? Evidence form Stock-Index Options," *Journal of Derivatives*, 7, 83–91.

Jackwerth, J. (1999): "Option Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review," *Journal of Derivatives*, 7, 66–82.

——— (2000): "Recovering Risk Aversion from Option Prices and Realized Returns," *Review of Financial Studies*, 13, 433–451.

——— (2004): *Option-Implied Risk-Neutral Distributions and Risk Aversion*, Charlotteville: Research Foundation of AIMR.

Jiang, G. and Y. Tian (2005): "The Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies*, 18, 1305–1342.

——— (2007): "Extracting Model-Free Volatility from Option Prices: An Examination of the VIX Index," *Journal of Derivatives*, 14, 1–26.

Kostakis, A., N. Panigirtzoglou, and G. Skiadopoulos (2011): "Market Timing with Option-Implied Distributions: A Forward-Looking Approach," *Management Science*, 57, 1231–1249.

Kumar, M., U. Moorthy, and W. Perraudin (2003): "Predicting Emerging Market Currency Crashes," *Journal of Empirical Finance*, 10, 427–454.

Li, H. and F. Zhao (2009): "Nonparametric Estimation of State-Price Densities Implicit in Interest Rate Cap Prices," *Review of Financial Studies*, 22, 4335–4376.

Low, C. (2004): "The Fear and Exuberance from Implied Volatility of S&P 100 Index Options," *Journal of Business*, 77, 527–546.

Lynch, D. and N. Panigirtzoglou (2008): "Summary Statistics of Option-Implied Probability Density Functions and Their Properties," Working Paper.

Malz, A. (1997): "Estimating the Probability Distribution of the Future Exchange Rate," *Journal of Derivatives*, 5, 18–36.

McMillan, L. (2004): *McMillan on Options*, New York: Wiley.

Melick, W. and C. Thomas (1997): "Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis," *Journal of Financial and Quantitative Analysis*, 32, 91–115.

Neumann, A. and G. Skiadopoulos (2012): "Predictable Dynamics in Higher Order Risk-Neutral Moments: Evidence from the S&P 500 Options," *Journal of Financial and Quantitative Analysis*, forthcoming.

Newey, W. and K. West (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Pan, X. (2012): "Investor Beliefs and State Price Densities in the Crude Oil Market," Working Paper.

Rosenberg, J. and R. Engle (2002): "Empirical Pricing Kernels," *Journal of Financial Economics*, 64, 341–372.

Schneider, P. (2012): "Fear Trading," Working Paper.

Schwert, W. (2011): "Stock Volatility during the Recent Financial Crisis," *European Financial Management*, 17, 789–805.

Whaley, R. (2000): "The Investor Fear Gauge," *Journal of Portfolio Management*, 26, 12–17.

——— (2009): "Understanding the VIX," *Journal of Portfolio Management*, 35, 98–105.

Moneyness		Call Options Days to Expiration			Put Options Days to Expiration		
		< 45	$45 - 90$	> 90	< 45	$45 - 90$	> 90
≤ 0.75	IV	103.03	64.29	50.29	88.02	66.04	53.55
	Obs.	5055	6213	10028	1507	3685	8552
	Avg Vol.	263	96	35	2566	786	260
	Avg OI	8320	3368	1205	32591	17487	5315
$0.75 - 0.95$	IV	74.66	64.63	54.12	77.07	65.22	54.66
	Obs.	5425	6648	12236	4961	6556	11986
	Avg Vol.	1529	551	153	5068	1720	342
	Avg OI	17151	9300	3223	41134	19882	5598
$0.95 - 1.05$	IV	84.71	71.91	58.46	84.43	71.95	58.88
	Obs.	2560	2779	5014	2560	2727	4745
	Avg Vol.	5305	1746	319	4489	1294	251
	Avg OI	36538	18846	4604	42330	15647	3874
$1.05 - 1.25$	IV	96.60	77.34	61.54	96.20	77.38	62.43
	Obs.	4366	4817	9172	4345	4556	7695
	Avg Vol.	6196	2121	284	1129	322	57
	Avg OI	50020	24960	5477	23201	7611	1299
> 1.25	IV	120.25	92.37	72.59	130.81	91.44	69.73
	Obs.	11207	16869	30878	14160	13482	17944
	Avg Vol.	3602	1245	211	73	19	8
	Avg OI	43321	17483	3650	2722	884	229

Table 1: VIX Options Data

The table shows implied volatilities (in percent) using the Black (1976) formula, the number of observations, the average volume, and the average open interest for different moneyness and maturity categories. Moneyness is defined as the strike price divided by the VIX futures price. The sample period is from February 24, 2006 to December 31, 2011.

	Date	VIX	VIX Return	SPX	SPX Return
1	February 27, 2007	18.31	64.22	1399.04	-3.47
2	August 8, 2011	48.00	50.00	1119.46	-6.66
3	August 4, 2011	31.66	35.41	1200.07	-4.78
4	August 18, 2011	42.67	35.12	1140.65	-4.46
5	September 29, 2008	46.72	34.48	1106.42	-8.81
6	May 6, 2010	32.80	31.67	1128.15	-3.24
7	November 9, 2011	36.16	31.59	1229.10	-3.67
8	October 22, 2008	69.65	31.14	896.78	-6.10
9	May 30, 2006	18.66	30.86	1259.87	-1.58
10	April 27, 2010	22.81	30.57	1183.71	-2.34
11	May 20, 2010	45.79	29.64	1071.59	-3.90
12	March 13, 2007	18.13	29.59	1377.95	-2.04
13	February 22, 2011	20.80	26.60	1315.44	-2.05
14	June 6, 2008	23.56	26.46	1360.68	-3.09
15	October 15, 2008	69.25	25.61	907.84	-9.03
16	November 1, 2007	23.21	25.26	1508.44	-2.64
17	May 7, 2010	40.95	24.85	1110.88	-1.53
18	October 19, 2007	22.96	24.11	1500.63	-2.56
19	January 28, 2011	20.04	24.09	1276.34	-1.79
20	October 30, 2009	30.69	23.95	1036.19	-2.81
21	December 1, 2008	68.51	23.93	816.21	-8.93
22	November 7, 2007	26.49	23.84	1475.62	-2.94
23	September 15, 2008	31.70	23.54	1192.70	-4.71
24	August 9, 2007	26.48	23.45	1453.09	-2.96
25	January 20, 2009	56.65	22.86	805.22	-5.28
26	July 13, 2006	17.79	22.77	1242.28	-1.30
27	January 22, 2010	27.31	22.63	1091.76	-2.21
28	August 10, 2011	42.99	22.62	1120.76	-4.42
29	October 31, 2011	29.96	22.14	1253.30	-2.47
30	May 17, 2006	16.26	21.80	1270.32	-1.68
31	January 20, 2006	14.56	21.54	1261.49	-1.83
32	March 16, 2011	29.40	20.89	1256.88	-1.95
33	November 27, 2009	24.74	20.80	1091.49	-1.72
34	February 4, 2010	26.08	20.74	1063.11	-3.11
35	June 4, 2010	35.48	20.43	1064.88	-3.44

Table 2: Largest Daily Percentage Changes

The table shows the largest daily percentage increases in the VIX, the levels of the VIX, the levels of the S&P 500 index (SPX), and the daily percentage changes in this index. The sample period is from February 24, 2006 to December 31, 2011.

	τ	N	Mean	Min	Max	Std. Dev.	Skewness	Kurtosis
Panel A: Levels								
RN Mean	30	1122	24.307	11.096	65.635	9.442	1.241	4.831
RN Mean	60	1463	24.084	11.814	60.657	8.780	1.016	4.305
RN Mean	120	1453	24.466	12.899	51.579	7.662	0.594	3.259
RN Variance	30	1122	38.797	2.354	400.316	48.060	3.285	16.704
RN Variance	60	1463	59.985	3.584	470.367	62.828	2.699	12.034
RN Variance	120	1453	82.644	10.260	426.456	65.871	1.931	7.115
RN Skewness	30	1122	1.271	0.695	1.933	0.261	0.209	2.215
RN Skewness	60	1463	1.413	0.832	2.074	0.254	0.189	2.293
RN Skewness	120	1453	1.485	0.935	2.198	0.243	-0.033	2.603
RN Kurtosis	30	1122	4.921	3.331	7.710	0.856	0.547	2.783
RN Kurtosis	60	1463	5.377	3.524	7.887	0.880	0.384	2.736
RN Kurtosis	120	1453	5.584	3.855	8.353	0.863	0.145	2.754
Panel B: First Differences								
RN Mean	30	1121	0.012	-6.576	8.880	1.276	0.666	9.631
RN Mean	60	1462	0.010	-5.051	5.577	0.953	0.503	8.657
RN Mean	120	1452	0.010	-2.929	4.438	0.671	0.535	8.461
RN Variance	30	1121	0.031	-147.511	138.607	16.405	0.064	32.029
RN Variance	60	1462	0.044	-135.114	134.087	15.431	0.342	27.381
RN Variance	120	1452	0.068	-60.206	91.238	9.815	1.078	19.052
RN Skewness	30	1121	0.000	-0.465	0.426	0.092	0.288	6.209
RN Skewness	60	1462	0.000	-0.440	0.558	0.066	0.598	10.421
RN Skewness	120	1452	0.000	-0.273	0.295	0.052	0.264	7.564
RN Kurtosis	30	1121	0.000	-2.121	1.594	0.355	-0.257	6.909
RN Kurtosis	60	1462	0.000	-1.783	1.734	0.264	0.230	9.255
RN Kurtosis	120	1452	0.001	-1.390	1.503	0.200	0.450	11.703

Table 3: Summary Statistics

The table shows descriptive statistics of the risk-neutral (RN) moments in levels (Panel A) and first differences (Panel B). τ denotes the time horizon of the option-implied moments in days and N the sample size. The sample period is from February 24, 2006 to December 31, 2011.

	RN Mean	RN Variance	RN Skewness	RN Kurtosis
Panel A: 30 days				
RN Mean	1.00	–	–	–
RN Variance	0.83	1.00	–	–
RN Skewness	-0.50	-0.22	1.00	–
RN Kurtosis	-0.49	-0.24	0.96	1.00
Panel B: 60 days				
RN Mean	1.00	–	–	–
RN Variance	0.86	1.00	–	–
RN Skewness	-0.40	-0.17	1.00	–
RN Kurtosis	-0.37	-0.17	0.97	1.00
Panel C: 120 days				
RN Mean	1.00	–	–	–
RN Variance	0.88	1.00	–	–
RN Skewness	-0.36	-0.15	1.00	–
RN Kurtosis	-0.27	-0.09	0.98	1.00

Table 4: Cross-Correlations

The table shows correlation coefficients between the risk-neutral (RN) moments over the sample period from February 24, 2006 to December 31, 2011. Panel A displays results for 30 days, Panel B for 60 days, and Panel C for 120 days.

	1	2	3	4	5	10	20	30	40	50	p
Panel A: Levels											
30 days											
RN Mean	0.99	0.98	0.97	0.96	0.96	0.92	0.83	0.73	0.64	0.58	0.18
RN Variance	0.94	0.91	0.89	0.86	0.84	0.74	0.53	0.38	0.29	0.26	0.00
RN Skewness	0.94	0.90	0.87	0.83	0.81	0.72	0.60	0.53	0.50	0.51	0.00
RN Kurtosis	0.91	0.87	0.83	0.78	0.76	0.68	0.56	0.51	0.48	0.47	0.00
60 days											
RN Mean	0.99	0.99	0.98	0.98	0.97	0.95	0.89	0.82	0.76	0.69	0.25
RN Variance	0.97	0.95	0.93	0.91	0.90	0.85	0.70	0.58	0.48	0.40	0.00
RN Skewness	0.97	0.95	0.93	0.91	0.90	0.83	0.74	0.64	0.59	0.56	0.00
RN Kurtosis	0.96	0.93	0.91	0.89	0.87	0.81	0.71	0.64	0.61	0.58	0.00
120 days											
RN Mean	1.00	0.99	0.99	0.98	0.98	0.96	0.90	0.85	0.79	0.73	0.33
RN Variance	0.99	0.98	0.96	0.95	0.94	0.89	0.76	0.66	0.57	0.48	0.04
RN Skewness	0.98	0.97	0.96	0.95	0.94	0.89	0.78	0.70	0.66	0.62	0.01
RN Kurtosis	0.97	0.96	0.95	0.94	0.93	0.88	0.77	0.70	0.66	0.62	0.01
Panel B: First Differences											
30 days											
RN Mean	-0.00	-0.10	0.00	-0.01	-0.08	0.05	-0.00	0.00	-0.05	0.04	0.00
RN Variance	-0.27	-0.04	0.04	-0.03	-0.07	-0.06	-0.01	-0.00	-0.02	0.03	0.00
RN Skewness	-0.23	0.02	-0.01	-0.06	-0.01	0.04	0.01	-0.01	-0.01	0.02	0.00
RN Kurtosis	-0.24	-0.04	0.03	-0.09	-0.02	0.06	0.04	-0.01	-0.02	-0.00	0.00
60 days											
RN Mean	-0.03	-0.04	0.00	-0.04	-0.05	0.10	0.04	0.00	0.02	-0.03	0.00
RN Variance	-0.16	-0.04	-0.00	-0.06	-0.08	0.12	0.01	0.02	-0.01	-0.04	0.00
RN Skewness	-0.25	0.04	-0.04	-0.05	0.06	0.02	-0.03	-0.03	0.00	-0.07	0.00
RN Kurtosis	-0.24	0.00	-0.02	-0.08	0.06	0.04	-0.01	0.00	-0.01	-0.10	0.00
120 days											
RN Mean	0.01	-0.02	-0.01	-0.00	-0.05	0.10	0.01	-0.00	0.01	-0.02	0.00
RN Variance	0.05	-0.01	-0.02	-0.02	-0.03	0.15	0.00	0.01	0.03	-0.03	0.00
RN Skewness	-0.29	0.02	-0.04	0.00	0.01	0.03	-0.01	0.00	0.01	0.01	0.00
RN Kurtosis	-0.32	0.05	-0.08	0.05	-0.00	-0.03	-0.02	0.01	0.01	0.03	0.00

Table 5: Autocorrelations

The table shows autocorrelations of the risk-neutral (RN) moments in levels (Panel A) and first differences (Panel B) over the sample period from February 24, 2006 to December 31, 2011. p refers to the p-value of the Phillips-Perron unit root test.

Quantiles	0.01	0.05	0.10	0.25	0.50	0.75	0.90	0.95	0.99
Panel A: 30 days									
Coefficient	0.11	0.16	0.20	0.29	0.41	0.63	0.94	1.10	1.48
t-statistic	4.85	6.87	8.33	11.37	16.90	27.61	23.05	21.12	17.17
Panel B: 60 days									
Coefficient	0.06	0.10	0.12	0.19	0.29	0.47	0.73	0.86	1.21
t-statistic	5.58	7.78	9.08	11.42	14.45	16.70	13.42	12.82	11.87
Panel C: 120 days									
Coefficient	0.06	0.07	0.09	0.14	0.21	0.31	0.44	0.51	0.70
t-statistic	9.52	10.88	11.41	12.24	12.31	11.49	10.04	9.75	9.22

Table 6: Relation Between VIX and Risk-Neutral Distribution

The table shows slope coefficients and t-statistics for regressions of the first-differenced risk-neutral quantiles onto the VIX first differences. The t-statistics are computed using the Newey and West (1987) adjustment with optimal lag length. Panel A displays results for 30 days, Panel B for 60 days, and Panel C for 120 days. The sample period is from February 24, 2006 to December 31, 2011.

	τ	VIX	RN Mean	RN Variance	RN Skewness	RN Kurtosis
Shanghai Surprise						
January 28, 2007	30	10.55	11.42	3.84	1.57	5.92
to	60	10.55	12.25	9.62	1.55	5.62
February 26, 2007	120	10.55	13.51	21.73	1.51	5.37
February 28, 2007	30	15.20	13.84	15.74	1.38	5.06
to	60	15.20	13.97	20.17	1.44	5.33
March 29, 2007	120	15.20	14.02	23.36	1.52	5.58
BNP Paribas						
July 10, 2007	30	19.21	18.54	22.34	1.30	4.93
to	60	19.21	18.27	31.36	1.39	5.24
August 8, 2007	120	19.21	18.40	38.86	1.32	4.98
August 10, 2007	30	25.54	23.23	55.81	1.40	4.81
to	60	25.54	22.54	83.80	1.55	5.55
September 8, 2007	120	25.54	21.32	86.66	1.62	5.95
Lehman Brothers						
August 16, 2008	30	21.88	22.69	18.25	1.01	4.03
to	60	21.88	22.93	31.53	1.11	4.27
September 14, 2008	120	21.88	23.27	45.40	1.11	4.23
September 16, 2008	30	45.31	31.79	100.15	1.13	4.32
to	60	45.31	29.88	100.44	1.18	4.52
October 15, 2008	120	45.31	27.43	75.54	1.03	4.14
TARP						
August 30, 2008	30	28.89	24.57	26.22	1.05	4.14
to	60	28.89	24.50	40.56	1.12	4.33
September 28, 2008	120	28.89	24.37	51.14	1.07	4.16
September 30, 2008	30	60.17	41.41	210.75	1.15	4.35
to	60	60.17	37.29	209.75	1.26	4.80
October 29, 2008	120	60.17	32.37	150.49	1.13	4.49
Flash Crash						
April 6, 2010	30	18.20	19.60	20.19	1.42	5.49
to	60	18.20	20.84	37.24	1.56	5.95
May 5, 2010	120	18.20	22.88	59.44	1.55	5.90
May 7, 2010	30	33.38	30.51	111.46	1.61	6.00
to	60	33.38	30.13	153.91	1.77	6.72
June 5, 2010	120	33.38	30.44	177.27	1.70	6.43
Downgrade						
July 6, 2011	30	20.73	19.63	25.27	1.66	6.27
to	60	20.73	20.20	41.08	1.85	7.00
August 4, 2011	120	20.73	21.34	59.04	1.83	6.82
August 6, 2011	30	36.80	31.55	98.54	1.46	5.29
to	60	36.80	28.90	115.53	1.75	6.50
September 4, 2011	120	36.80	27.23	114.13	1.85	6.99

Table 7: Summary Statistics for Crash Periods

The table shows descriptive statistics of the VIX and the risk-neutral (RN) moments 30 days before and after each event. τ denotes the time horizon in days.

	Coefficient	s.e.	p	Mean $Y=1$	Mean $Y=0$
Panel A: 97.5% Quantile					
RN Variance	18.9788	5.1295	0.0002	0.0495	0.0407
RN Skewness	0.7859	0.5556	0.1572	0.6551	0.6366
Put/Call (OI)	-0.0999	0.7284	0.8910	0.5244	0.4942
Put/Call (Volume)	0.0176	0.0781	0.8214	0.5890	0.5759
Yield Spread	0.1495	0.1178	0.2043	2.2588	1.9364
Constant	-2.1060	0.6616	0.0015		
$\ln L$	-653.1641				
$\ln L_0$	-696.3653				
R^2	0.0620				
Panel B: 99% Quantile					
RN Variance	17.0284	5.2760	0.0012	0.0516	0.0421
RN Skewness	1.3103	0.6628	0.0480	0.7071	0.6319
Put/Call (OI)	-0.1279	1.0080	0.8990	0.5265	0.4999
Put/Call (Volume)	-0.0174	0.1002	0.8620	0.5658	0.5823
Yield spread	0.1401	0.1436	0.3293	2.3006	1.9944
Constant	-2.9506	0.8430	0.0005		
$\ln L$	-416.3279				
$\ln L_0$	-452.4042				
R^2	0.0797				

Table 8: Probit Model

The table shows the parameter estimates of the probit model. The explanatory variables are: 1. risk-neutral (RN) variance, 2. risk-neutral (RN) skewness, 3. put to call ratio based on open interest (OI), 4. put to call ratio based on volume, and 5. yield spread. The standard errors (s.e.) of the coefficient estimates are computed using the Newey and West (1987) estimator with optimal lag length. p refers to the p-value of the null hypothesis that the true coefficient value is zero. $\ln L$ is the maximized value of the log-likelihood function, $\ln L_0$ the restricted (with only a constant term) log-likelihood, and R^2 denotes McFadden's pseudo R-squared. The two columns on the right show the average values of the independent variables given an upward spike ($Y = 1$) and no jump ($Y = 0$). The sample period is from February 24, 2006 to December 31, 2011.

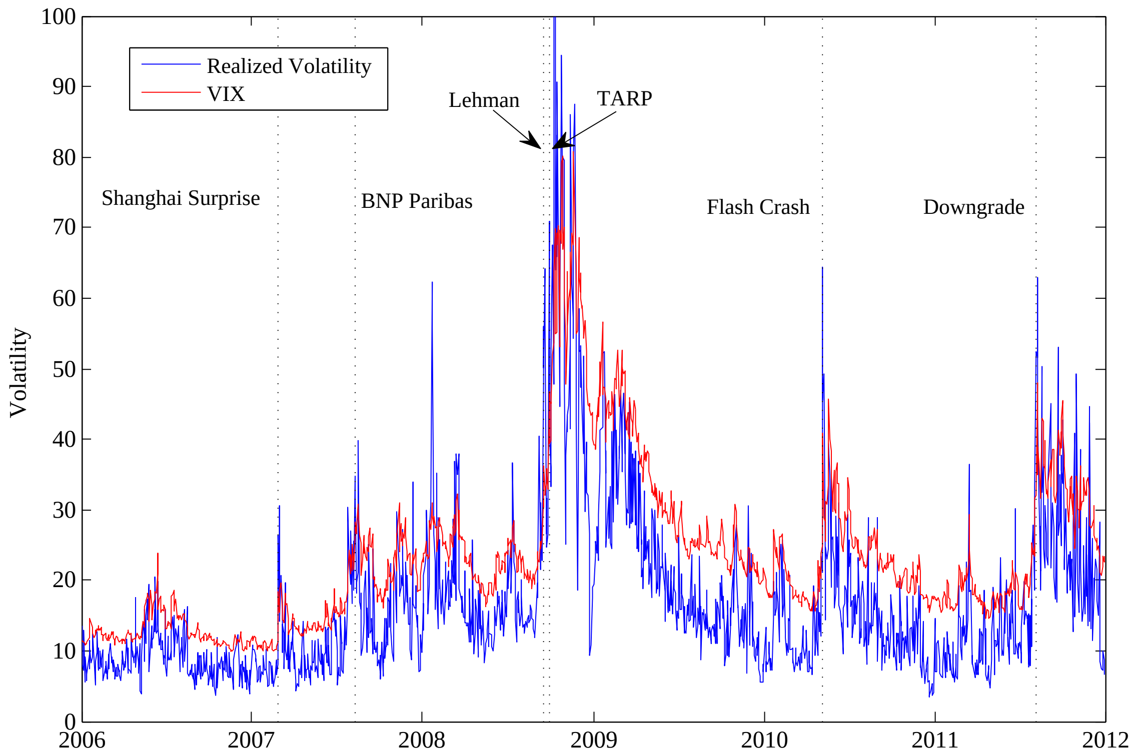

Figure 1: Realized Volatility and VIX

The figure shows realized volatilities and the VIX. Realized volatilities are calculated by summing up squared 5-minute intra-day returns and afterwards taking the square-root. The sample period is from February 24, 2006 to December 31, 2011.

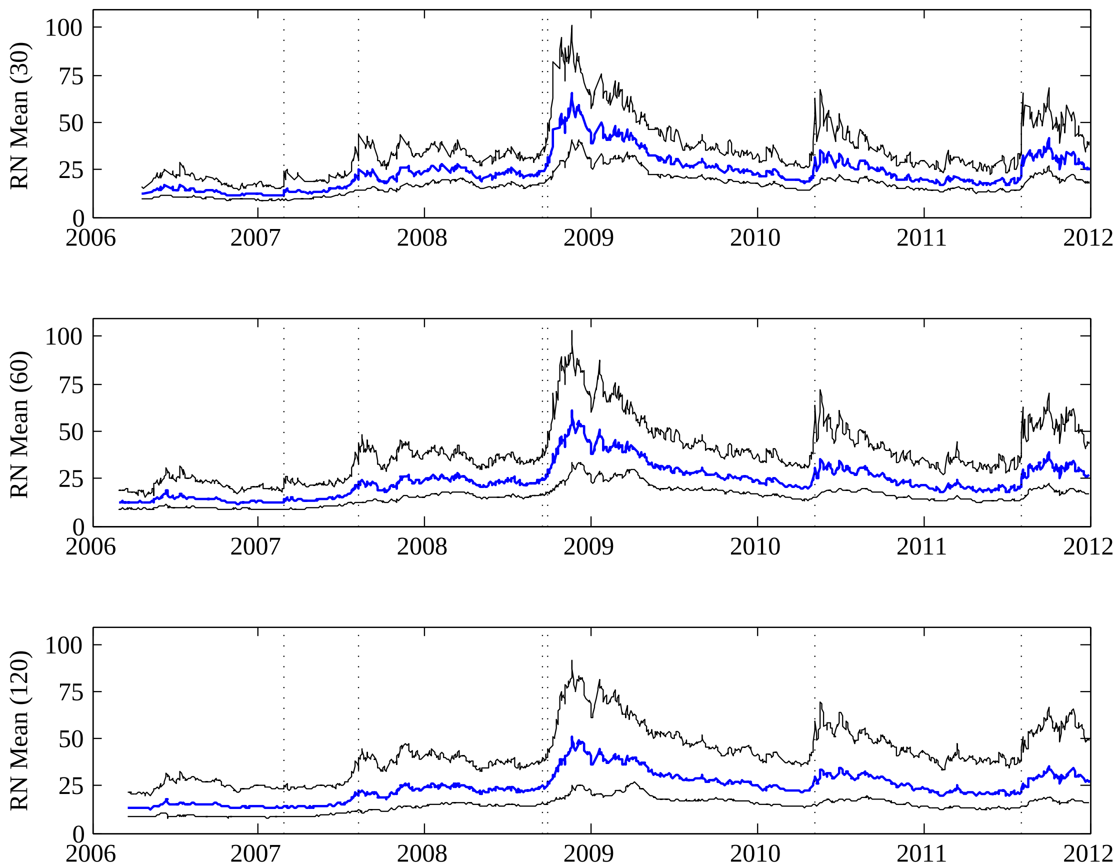

Figure 2: Risk-Neutral Mean

The figure shows time-series of the risk-neutral (RN) mean (blue) and the 90% confidence interval (black) for three times to maturity (30, 60, and 120 days). The dotted vertical lines correspond to the same events as in Figure 1. The sample period is from February 24, 2006 to December 31, 2011.

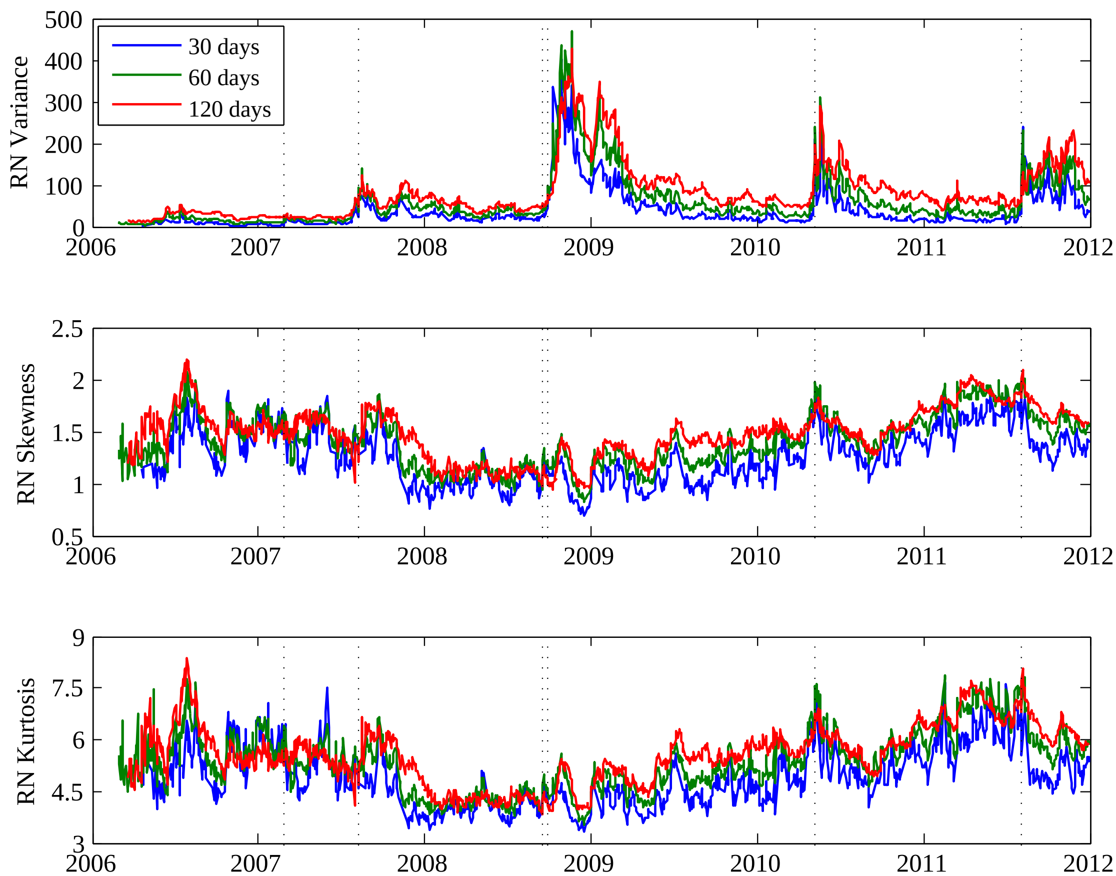

Figure 3: Risk-Neutral Moments

The figure shows time-series of the variance, skewness, and kurtosis of the risk-neutral (RN) distribution for three times to maturity (30, 60, and 120 days). The dotted vertical lines correspond to the same events as in Figure 1. The sample period is from February 24, 2006 to December 31, 2011.

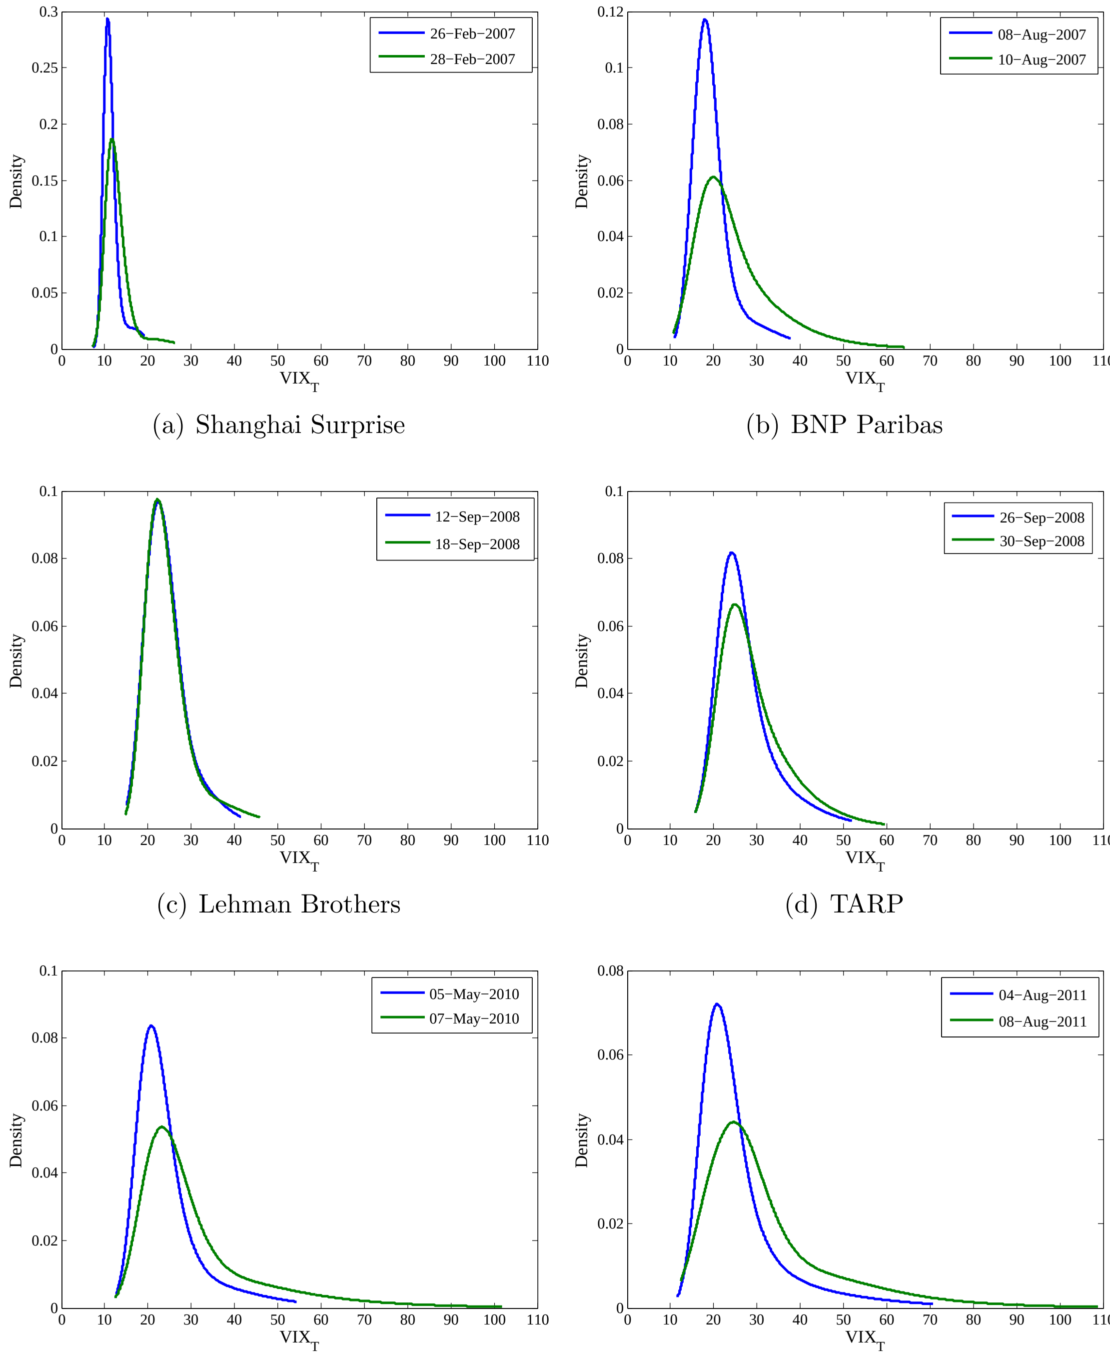

(a) Shanghai Surprise (b) BNP Paribas

(c) Lehman Brothers (d) TARP

(e) Flash Crash (f) Downgrade

Figure 4: Risk-Neutral Distribution around Selected Events

The figure shows risk-neutral distributions one day before and after each event. The time horizon is set to 30 days.

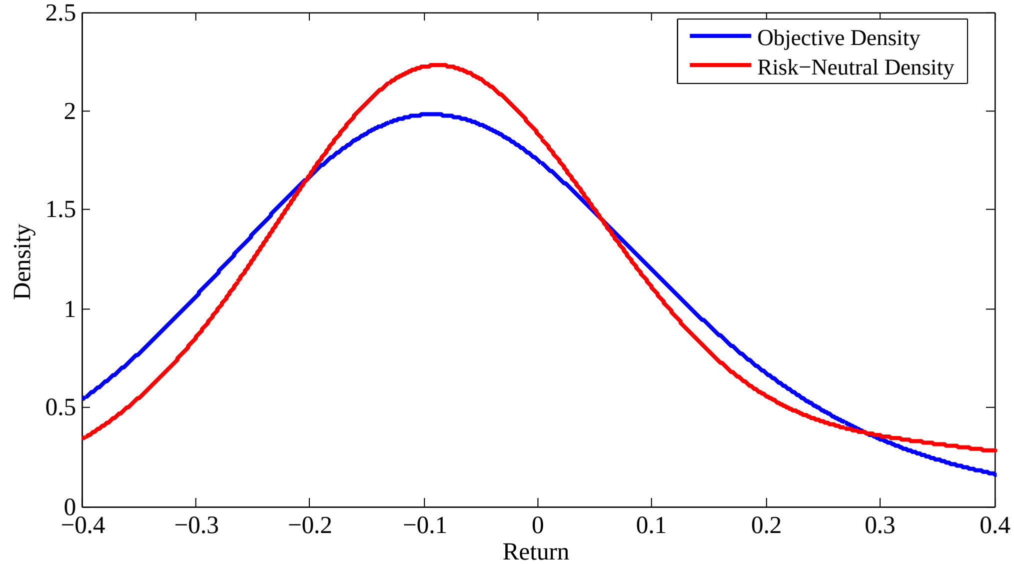

(a) Densities

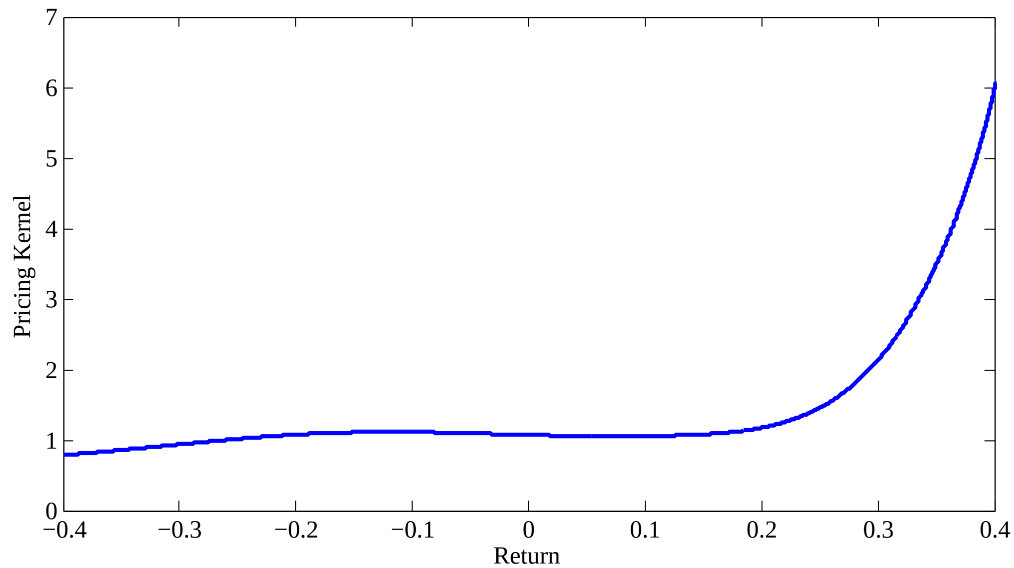

(b) Pricing Kernel

Figure 5: Densities and Volatility Pricing Kernel

The figure shows the average risk-neutral distribution, objective distribution, and volatility pricing kernel. The densities and the pricing kernel are expressed in terms of log returns. The time horizon is set to 30 days. The sample period is from February 24, 2006 to December 31, 2011.

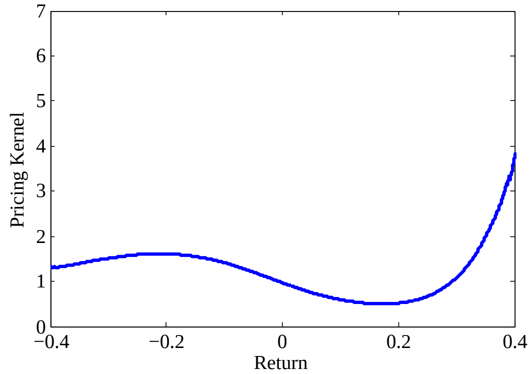

(a) 2006

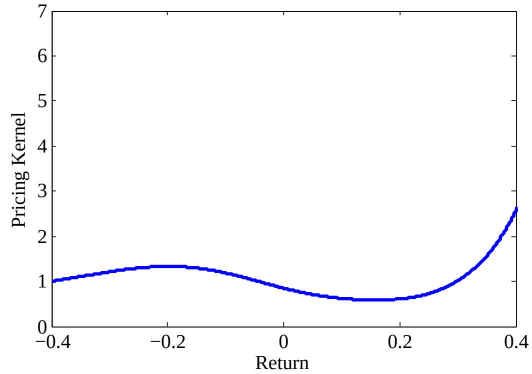

(b) 2007

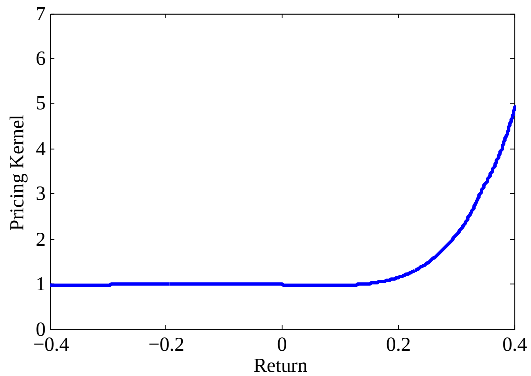

(c) 2008

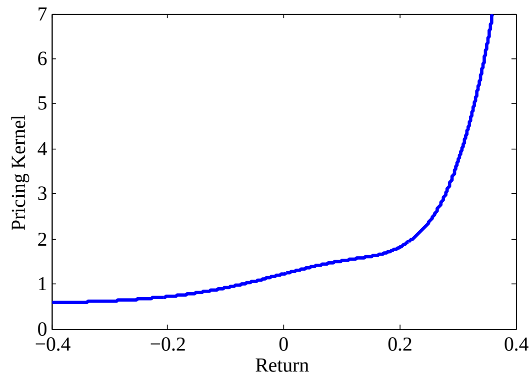

(d) 2009

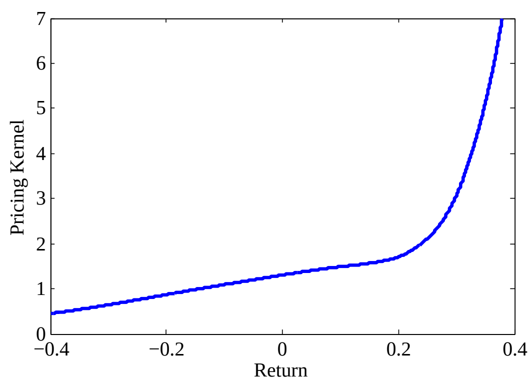

(e) 2010

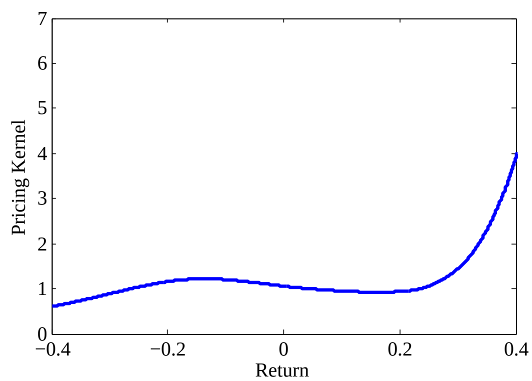

(f) 2011

Figure 6: Volatility Pricing Kernel (Yearly Averages)

The figure shows the average volatility pricing kernel expressed in terms of log returns for each year of the sample period from February 24, 2006 to December 31, 2011. The time horizon is set to 30 days.