# Uncovering Novel Features of Equity-Index Return Dynamics via Corridor Implied Volatility*

Torben G. Andersen† Oleg Bondarenko‡ Maria T. Gonzalez-Perez§

July 2010; Revised: September 2012

**Abstract**

A number of fundamental questions regarding the equity-index return dynamics are difficult to address due to the latent character of spot volatility. For example, it is hard to discern much about the nature of volatility jumps or the features of the negative (spot) return-volatility correlation. We show that, in theory, the VIX index can be useful in addressing such issues. Unfortunately, in practice, the high-frequency VIX series is plagued by idiosyncratic biases and noise which severely distort the measure. Instead, we exploit high-frequency option quotes to compute a novel "Corridor Implied Volatility" index (CX) with superior properties. Exploiting the CX index, we obtain striking new empirical findings regarding the nature of volatility jumps and the return-volatility asymmetries. Finally, we document that the CX index is robust and performs admirably, even during turbulent market conditions. In contrast, the existing VIX measure is particularly unreliable during periods of market stress, exactly when a "fear gauge" is most valuable.

*JEL* Classification: G13, C58

*Keywords*: VIX, Model-Free Implied Volatility, Corridor Implied Volatility, Time Series Coherence

*Elements of this work were previously included in the working paper entitled "Coherent Model-Free Impled Volatility: A Corridor Fix for High-Frequency VIX." We are indebted to the Zell Center for Risk at the Kellogg School of Management, Northwestern University, for financial support. We thank Kris Jacobs, Jesper Lund and seminar participants at the European Finance Association Meeting, Copenhagen, August 2012, the 2012 Western Finance Association Meeting, Las Vegas, the 2011 Osaka Conference on "High-Frequency Data Analysis in Financial Markets," the Midwest Econometrics Group 2011 Meetings at the University of Chicago, the Federal Reserve Board of Governors, the SoFiE-CREATES 2010 Conference in Aarhus, the University of Illinois at Chicago, Virginia Tech, Spot Trading, the Wisconsin School of Business, and the Illinois Institute of Technology for comments.

†Kellogg School of Management; Northwestern University; 2001 Sheridan Road, Evanston, IL 60208, NBER, and CREATES; e-mail: t-andersen@kellogg.northwestern.edu

‡Department of Finance (MC 168), University of Illinois at Chicago, 601 S. Morgan St., Chicago, IL 60607; e-mail: olegb@uic.edu

§Colegio Universitario de Estudios Financieros (CUNEF); Serrano Anguita, 9; 28004 Madrid, Spain; e-mail: m-gonzalezperez@cunef.edu

# 1 Introduction

The characterization of the equity-index return dynamics has been refined in recent decades with the growing access to daily, and subsequently tick-by-tick, financial price data; yet much remains unresolved. We now recognize that returns have time-varying means, that they display very persistent volatility fluctuations, that market prices occasionally move abruptly, i.e., "jump," and that return and volatility innovations are negatively correlated. We are coming to grips with the fact that return volatilities jump as well and, indeed, volatility and returns may jump simultaneously, i.e., "co-jump." Nevertheless, it remains elusive how these effects originate, interact and propagate from the highest intraday frequencies to generate the observed fat-tailed and asymmetric return distributions at longer horizons.

On the surface, this may appear to be a peripheral issue. If the real economy responds to events at a lower, say monthly, frequency, we can use monthly returns to construct models for the conditional return dynamics at corresponding and lower frequencies. However, it is widely recognized that daily returns provide superior inputs for estimating the conditional return distribution. The reason is twofold. First, return volatility is unobserved, or latent, so the current state of volatility cannot be identified with certainty. Using monthly data, the update of volatility from one month to the next is based on a single observation. It is impossible to discern whether volatility has been trending up or down prior to month end. Given the strong persistence in volatility, accurate assessment of the current state is critical in forming useful predictions about future price fluctuations. Daily data clearly facilitate improved inference in this regard. Second, the monthly sampling frequency is too sparse when the environment is turbulent. Policy makers, business leaders and financial market participants must interpret, and respond to, ongoing developments in real time, as economic and financial shocks hit the markets. Under such circumstances, swift updates to the estimate of market conditions is of utmost importance.

In principle, the identical argument justifies the move from daily to intraday data. In fact, this is arguably key to proper measurement of critical features of the return generating process, including the size and prevalence of jumps and the origin of the dynamic interaction between return and volatility. However, new challenges arise. The volatility process is subject to pronounced intraday patterns and short-lived bursts associated with real-time price discovery in response to news arrivals. Moreover, institutional features and market microstructure effects severely reduce the signal-to-noise ratio of individual returns at the highest frequencies. The widely adopted realized volatility (RV) approach "solves" this problem by integrating the volatility over the trading day. That is, the intraday squared returns at modestly high intraday frequencies are aggregated to provide an estimate of the daily return variation, thus harnessing the information content of intraday price fluctuations while avoiding the delicate task of modeling intraday patterns, news effects, and microstructure frictions.[1] However, this precludes a deeper inquiry into the high-frequency return dynamics. The daily RV measure reflects the average volatility. It is silent regarding the state of volatility at the close of trading, and it does not speak directly to the prevalence of price and volatility jumps, or their interaction.[2]

The development of reliable tools for estimation of the (latent) high-frequency evolution of *spot* volatility, including the timing and size of volatility jumps, has proven even more challenging. The problem arises from two offsetting effects. One, in order to accommodate microstructure frictions, volatility must be extracted from returns that each cover a non-trivial time interval. Next, these individual returns must be combined over a longer time interval to assure a sufficient number of observations for efficient volatility measurement. As a consequence, these procedures, effectively, rely on volatility being (near) constant over sizeable intraday periods. Two, the intraday volatility pattern, volatility

[1] See, e.g., Andersen, Bollerslev & Diebold (2010) for a general review.

[2] Subsequently, several tests for the presence of price jumps have been devised, see, e.g., Barndorff-Nielsen & Shephard (2004), Huang & Tauchen (2005), Lee & Mykland (2007), Andersen, Bollerslev & Dobrev (2007), and Mancini (2009).

spikes associated with news arrivals, as well as price and volatility jumps limit the length of the the intraday intervals over which volatility can reasonably be assumed to remain invariant. Ultimately, there is no perfect solution. Ait-Sahalia, Fan & Li (2012) also argue, forcefully, that these estimation procedures for spot volatility cannot provide a viable basis for inference about the leverage effect.

In this paper, we propose exploiting high-frequency observations on model-free implied volatility (MFIV) measures derived from options to analyze important characteristics of volatility jumps as well as the instantaneous return-volatility correlation. There are two critical issues to address before adopting this approach. First, we must justify that the option-implied measures can convey meaningful information about spot volatility. The MFIV measures reflect the longer horizon (risk-neutral) volatility expectations embedded in option prices. Nonetheless, under general and commonly adopted assumptions, the risk-neutral volatility dynamics is governed by a smooth function of the state vector which includes (components of) spot volatility, implying that the MFIV measures themselves typically are smooth, monotonically increasing functions of actual spot volatility. In particular, the MFIV measure jumps if spot volatility jumps. Second, we must verify that we can identify the spot return-volatility correlation from high-frequency correlations between returns and MFIV measures. Again, this follows under standard assumptions and, in particular, the equivalence of the risk-neutral and the actual (or physical) probability measures. In practice, the crucial requirement for accomplishing these tasks is the availability of reliable high-frequency volatility index observations.

The most natural choice of a high-frequency MFIV measure is the VIX, disseminated by the Chicago Board of Options Exchange (CBOE) about every 15 second during the trading day. As always, the main concern is the potential distortions arising from microstructure effects. We identify an influential source of error that is particularly detrimental for our applications, namely the presence of random and occasionally abrupt fluctuations in the range of options exploited in the VIX computations. Such sudden shifts induce artificial breaks or "jumps" in the index. Effectively, the VIX is capturing a different underlying notion of volatility depending on the strike range. We refer to this feature as *lack of coherence*. As a result, the volatility series often "jumps," not due to a change in volatility or option prices, but due to a shift in the range of options used to compute the measure. We conclude that the high-frequency VIX is not suited for our purposes.

Instead, we turn to an alternative methodology for volatility index construction. We deliberately focus on a limited strike range which, measured by a suitable option metric, covers an invariant portion of the risk-neutral density. Hence, we introduce a new cut-off criterion, determined endogenously by option prices, that allows the volatility measure to reflect an economically equivalent fraction of the strike range, ensuring *intertemporal coherence* of the measure. We label this measure a model-free corridor implied volatility index (CX). It is closely related to the corresponding notion discussed by Carr & Madan (1998). We calibrate our corridor to obtain good coverage of the strike range while ensuring the availability of reliable quotes for the relevant options across the entire sample.

We compute this CX index every 15-seconds throughout the regular trading day using quotes on the CBOE SPX options over a 25-month period which includes the extremely tumultuous market conditions associated with the financial crisis in the fall of 2008 and the flash crash on May 6, 2010. The index construction necessitates large-scale data cleaning as each CX computation, on average, exploits more than two hundred pairs of bid-ask option quotes across two separate maturities. Identifying all relevant option quotes and checking for various types of data error is crucial for the quality of our implied volatility index. To the best of our knowledge, it represents the first use of a comprehensive tick-by-tick option data set for addressing issues within this general area of research.

Our empirical findings, exploiting the high-frequency CX index, are striking. First, we document that large changes in volatility are common and the volatility jump distribution is close to symmetric.

In addition, volatility jumps are often accompanied by price jumps in the opposite direction. These conclusions are obtained from direct observations of the volatility index at the one-minute frequency and, as such, operate over short horizons for which traditional high-frequency methods, exploiting only return observations, cannot provide reliable inference. Our findings are at odds with prior empirical work both in terms of the jump frequency and jump size distribution for volatility. They also point towards misspecification of popular asset pricing models that, directly or indirectly, assume that volatility only jumps upward. Another noteworthy feature is the actual size of the jumps. The largest two one-day (absolute) CX moves are 35% and 30%, while the largest one-minute changes in CX were 7.13% (during the flash crash) and 3.89%. Similarly, the largest daily price moves were 13.2% and 11.8% versus one-minute changes of 2.05% and 2.01%. That is, the size of the discontinuities is an order of magnitude smaller than one may infer from daily data. Instead, the largest daily moves are generated by dynamic high-frequency interactions between the return and volatility innovations and, possibly, a few smaller jumps. Thus, while it is not surprising we uncover more jumps than studies based on daily data – we effectively exploit a refined "microscope" – the differences are economically significant.

Second, we find the spot return-volatility correlation – also known as the "leverage effect" – to be very pronounced, with the coefficient ranging between $-0.68$ and $-0.91$ over our sample period. These estimates are significantly more negative than what is reported in the extant literature. Moreover, we present evidence of a substantial time variation in the leverage effect which grows more extreme in times of economic uncertainty when the volatility index is elevated. This type of detailed investigation of systematic fluctuations in the leverage effect is again contingent on our ability to directly exploit an observable volatility index at a very high frequency. Moreover, we reiterate that these conclusions simply cannot be established using the high-frequency VIX series.

Finally, we illustrate how the replacement of the VIX index by the CX series is pivotal also for our ability to monitor real-time developments during turbulent market conditions. We document large persistent biases in the VIX measure during the critical phase of the so-called flash crash on May 6, 2010. In contrast, the CX index seems to accurately reflect the dynamics in a way that is coherent with the evolution of the S&P 500 futures traded concurrently at CME Group exchange. The failure of the VIX during the flash crash is symptomatic of the inherent problems with the index during volatile market conditions. In those exact scenarios where we most value precise and timely information regarding volatility movements and dynamic interactions among returns and volatility, the VIX measure becomes so distorted that meaningful real-time assessment is precluded. We argue that it would be useful to have exchanges compute a real-time volatility index for the S&P 500 along the lines of our CX measure to enhance transparency and monitoring of market conditions during period of market stress.[3]

The remainder of the paper proceeds as follows. Section 2 presents the methodology underlying the VIX computation and highlights the potential pitfalls. Section 3 reviews our data sources. Section 4 documents practical problems with the high-frequency VIX measure. The insights obtained here rationalize our construction of the corridor implied volatility measure in Section 5. Section 6 describes the empirical analysis in three separate subsections dealing with volatility jumps, the leverage effect, and the flash crash. Section 7 concludes, while the appendix contains additional information regarding computational and inferential features of our analysis.

[3]The VIX futures, traded on an electronic platform at the CBOE, are often seen as a suitable alternative, but trading in this market vanished for an extended period of time during the flash crash.

# 2 The Methodology behind the VIX

The CBOE volatility index, VIX, is designed to approximate the so-called *model-free implied volatility*, or MFIV. This is an option implied measure, based on theory demonstrating that, under quite general conditions, the "fair" value of a claim promising to pay the future realized asset return variation over a specific horizon is given by the market price of a particular portfolio of European out-of-the-money (OTM) call and put options. The CBOE computes and disseminates the value of this index almost continuously during the trading day, effectively rendering the VIX measure observable in real time. This section reviews the basic concepts as well as the practical implementation of the VIX calculation. It sets the stage for our subsequent development of a computational procedure that provides an alternative measure which, in practice, can be more effectively aligned with an invariant notion of volatility.

## 2.1 Theory and Implementation for the VIX Computation

The computation of the VIX is based on the notion of model-free option-implied volatility. This concept is most readily grasped within a purely diffusive setting, and we organize the exposition accordingly, although we introduce jumps later when they are pertinent to the issues under consideration.

We let the logarithmic equity-index value at time $t$ be denoted $s_t = log(S_t)$. In a frictionless market, if the price path is continuous and there are no arbitrage opportunities, the price evolves according to a Brownian semi-martingale, see, e.g., Back (1991),

$$ds_t \;=\; \mu_t\, dt \;+\; \nu_t\, dW_t\,, \tag{1}$$

where $W_t$ denotes a standard Brownian motion. Hence, the log price dynamics is characterized by a (stochastic volatility) diffusion. Under weak regularity conditions, and absent specific parametric assumptions regarding the drift and diffusion coefficients, $\mu_t$ and $\nu_t$, the quadratic return variation, or *integrated variance*, over the time interval $[0,T]$, denoted *IVar*, is given by,

$$IVar \;=\; \int_0^T \nu_u^2\, du\,.$$

The integrated variance is the natural notion of (realized) return variation in this setting, see, e.g., Andersen, Bollerslev & Diebold (2010). One way to estimate the expected future return variation is to construct a security with no intermediate cash flows and a terminal (random) payoff equal to *IVar*. This security would, in theory, command the price, $E^*[IVar]$, where $E^*[\cdot]$ denotes the expected value operator under the risk-neutral measure. It is readily converted into an annualized volatility measure,

$$\sigma_T^2 \;=\; \frac{1}{T}\, E^*\left[\, \int_0^T \nu_u^2\, du \,\right] \;=\; \frac{1}{T}\, E^*\left[\, IVar \,\right]. \tag{2}$$

In fact, a trading strategy with a payoff equal to the integrated variance can be constructed. It combines a static option portfolio with self-financing transactions in the underlying asset. Consequently, the price of the integrated variance (payoff) simply equals the cost of the option portfolio. Formally,

the valuation formula takes the form,[4]

$$\sigma_T^2 \;=\; \frac{2\,e^{rT}}{T}\left[\int_0^F \frac{P(T,K)}{K^2}\,dK + \int_F^\infty \frac{C(T,K)}{K^2}\,dK\right] \;=\; \frac{2e^{rT}}{T}\left[\int_0^\infty \frac{Q(T,K)}{K^2}\,dK\right], \tag{3}$$

where $r$ is the (annualized) risk-free interest rate for the period $[0,T]$, measured by the corresponding U.S. Treasury bill rate, $T$ is time-to-maturity measured in units of a year, $F$ denotes the forward price for maturity $T$, $P(T,K)$ and $C(T,K)$ are the prices for European put and call options with strike $K$ and time-to-maturity $T$, and $Q(T,K) = \min\{C(T,K), P(T,K)\}$ denotes the price of the OTM option at strike $K$. Hence, for $K < F\,(K > F)$, $Q(T,K)$ is the price of an OTM put (call). Correspondingly, we label an option with strike $K = F$, and thus $C(T,K) = P(T,K)$, an "at-the-money" (ATM) option.

Evidently, the computation of the model-free implied variance requires the availability of market prices for a continuum of European-style options with strike prices spanning the support of all possible future index levels, from zero to infinity, for the relevant maturity. In practice, the number of available strikes is limited so the CBOE approximates $\sigma_T^2$ in equation (3) by the following quantity,

$$\hat{\sigma}_T^2 \;=\; \underbrace{\frac{2\,e^{rT}}{T}\sum_{i=1}^{n}\frac{\Delta K_i}{K_i^2}\,Q(T,K_i)}_{\text{Discrete Approximation}} - \underbrace{\frac{1}{T}\left[\frac{F}{K_f}-1\right]^2}_{\text{Correction Term}}, \tag{4}$$

where $0 < K_1 < \cdots < K_f \leq F < K_{f+1} < \cdots < K_n$ refer to the strikes included in the computation, $K_f$ denotes the first strike price available below the forward rate, $F$, where the index $f$ is a positive integer, and we assume a reasonable cross-section of options are available so that, $2 < f < n-1$.[5] Finally, the second term in (4) reflects a correction for the discrepancy between $K_f$ and the forward price.

Another complication is that, for organized exchanges, only a few option maturity dates are quoted at any given time, while the $VIX$ is defined for a fixed calendar maturity of $T_M = \frac{30}{365}$, or thirty days. The CBOE obtains the $VIX$ measure by linearly combining $\sigma_1^2$ and $\sigma_2^2$ for the two expiration dates closest to thirty calendar days with time-to-maturity of $T_1$ and $T_2$, but excluding options with less than seven calendar days to expiry. The interpolated quantity is expressed as a volatility measure and quoted in terms of annualized percent,

$$VIX = 100 \times \sqrt{\left[w_1\left(\hat{\sigma}_1^2 T_1\right) + w_2\left(\hat{\sigma}_2^2 T_2\right)\right] \times \frac{365}{30}}, \tag{5}$$

where $w_1 = \frac{T_2 - T_M}{T_2 - T_1}$ and $w_2 = \frac{T_M - T_1}{T_2 - T_1}$, so that, obviously, $w_1 + w_2 = 1$.

In summary, all inputs for the VIX, apart from the risk-free rate $r$ and the forward rate $F$, are based on options data. Section 4.1 explains how the forward rate is derived via put-call parity. As a result, the VIX index is calculated from CBOE option quotes and U.S. Treasury yields exclusively.

[4] Carr & Wu (2009) and Bondarenko (2010) demonstrate that this formula provides a very good approximation to the price of the future quadratic return variation (the integrated variance *and* the realized squared jumps) – even in the presence of price jumps – as long as the jumps are not extremely large. Martin (2012) develops an alternative option-based procedure pricing the risk-neutral variance that accommodates jumps and does not involve any approximation error.

[5] The increments in the strike ranges are calculated as $\Delta K_1 = K_2 - K_1$, $\Delta K_n = K_n - K_{n-1}$, and for $1 < i < n$, $\Delta K_i = (K_{i+1} - K_{i-1})/2$, while the term $Q(T,K_i)$ is defined as the midpoint of the bid-ask spread for the OTM option with strike $K_i$. More specifically, $Q(T,K_i)$ equals the put price when $K_i < K_f$, the call price when $K_i > K_f$, and the average of the two when $K_i = K_f$. Additional information is available at http://www.cboe.com/micro/vix/vixwhite.pdf

## 2.2 Tail Truncation as a Source of Random Bias in the VIX

There are various sources of measurement error in the VIX, computed using equation (4). Jiang & Tian (2005) classify them as follows: (i) truncation errors – the minimum and maximum strikes are far from zero and infinity; (ii) discretization errors – piecewise linear functions approximate the integrals; (iii) "expansion" errors – a Taylor series expansion approximates the log function in deriving the correction term; and (iv) interpolation errors – linear interpolation of the maturities.

Our objective is to exploit high-frequency movements in the volatility index for studying important dynamic features of equity-index returns. In that context, we identify a qualitatively different source of error which dominates any of those above: it is critical to control for *idiosyncratic fluctuations* in the VIX, stemming from the procedure which determines the strike range used in computing $\hat{\sigma}^2$.

To appreciate the potential impact of this random truncation bias, it is necessary to understand the CBOE rule determining the strike range used for computing the VIX. In principle, all (available) options should be exploited. However, the relative bid-ask spread is very high for far OTM options and even midpoint quotes may be poor indicators of underlying value. The CBOE adapts to this feature by invoking a specific cut-off rule: moving away from the forward price, once two consecutive strikes with zero bid quotes are encountered, all further OTM options are discarded. Hence, market liquidity and pricing jointly determines the extreme strikes, $K_1$ and $K_n$, exploited in the index computation. As we show in the empirical section below, this cut-off rule for the OTM options produces a *time-varying* effective strike range which induces occasional spurious breaks in the VIX series, entirely unrelated with concurrent developments in the underlying option values or return series (*artificial volatility jumps*).

To quantify the implications of this time-varying truncation bias, we formalize the notion of measuring the expected return variation via option prices covering only a limited strike range. Specifically, we define the indicator function, $I_t(B_1, B_2)$, for barriers $B_1$ and $B_2$, $0 < B_1 < B_2 < \infty$, as,

$$I_t(B_1,B_2) \;=\; 1[B_1 \leq S_t \leq B_2]\,,$$

which takes the value of unity if the asset price at time $t$ is within $[B_1, B_2]$ and zero otherwise.

Carr & Madan (1998) demonstrate that equation (2) naturally generalizes to the case involving a finite strike range, as the measure then captures the so-called *Corridor Integrated Variance*, or CIVar,

$$\frac{1}{T}\,E^*\left[CIVar(B_1,B_2)\right] \;\equiv\; \frac{1}{T}\,E^*\left[\int_0^T v_u^2\, I_u(B_1,B_2)\,du\right] \;=\; \frac{2\,e^{rT}}{T}\int_{B_1}^{B_2}\frac{Q(T,K)}{K^2}\,dK\,. \tag{6}$$

In other words, the corridor integrated variance reflects the return variation that is realized only while the asset price is within prespecified barriers. Clearly, the corridor variance is a downward biased measure of MFIV, as it excludes the contribution from OTM options outside the corridor.

The key point is that the VIX index in equation (4) is akin to a randomly truncated variance measure,

$$\frac{1}{T}\,E^*\left[CIVar(K_1,K_n)\right] \;=\; \frac{2\,e^{rT}}{T}\int_{K_1}^{K_n}\frac{Q(T,K)}{K^2}\,dK\,, \tag{7}$$

with barriers determined by the CBOE truncation rule. Because the tails are not accounted for, VIX is downward biased. If this bias represents a fixed proportion of MFIV, the time variation in VIX still provides a (scaled) measure of the underlying fluctuations in model-free volatility. However, importantly, there is a random component to the bias, as the strike range is subject to idiosyncratic fluctuations, stemming from the application of the truncation rule to a market with shifting liquidity and quote intensity. If this effect is substantial, the VIX measure becomes *incoherent*: it does not capture an

invariant notion of volatility. It reflects not only option prices (volatility) but also random shifts in the width of the corridor. We document below that this feature is critical for empirical work seeking to relate high-frequency volatility innovations to the dynamic behavior of the underlying return process.

Of course, the lack of precise indicators of option value in the tails of the distribution is ubiquitous. It makes some kind of truncation inevitable in practice. We argue, later on, that the problem can be greatly alleviated via a novel implementation of the Corridor Implied Volatility index, initially developed in Andersen & Bondarenko (2007). This construction explicitly controls the range $(K_1, K_n)$ in a time-consistent manner to ensure a coherent basis for computing model-free volatility.

# 3 Data

We compute model-free volatility indices from SPX option quotes obtained from Market Data Express, henceforth MDX. The sample covers June 2, 2008 – June 30, 2010, encompassing 525 trading days. For the two relevant maturities, we have an average of about 213,000 (147,000) OTM put (call) option quotes per day.[6] We construct fifteen second series for each of these options using the "previous tick" method across the active trading day, from 8:30:15 to 15:15:00 Chicago Time (CT).[7] If no quote arrives in a 15 second interval, the last available quote prior to the interval is used. In the case of inactivity, the quotes may become stale. Hence, we impose a limit on the duration of widespread inactivity as described below. In addition, as a benchmark, we occasionally exploit the corresponding intraday quotes for the S&P 500 futures from the CME Group.

A few filters are applied to guard against data errors. First, we avoid excessive staleness by classifying an option as missing if the quote time-stamp precedes the volatility index by more than five minutes. If quotes for nearby strikes are available, the VIX computation interpolates the option value for this strike via the formula in equation (4), so this does not create a gap in the volatility series. Second, if an entire block of adjacent and relevant OTM option prices have not been updated for five minutes, we deem the index itself "not available" (n.a.) and eliminate the VIX observations across this entire period to avoid errors due to "systemic staleness." This condition is critical for avoiding artificial jumps when a set of stale quotes are updated simultaneously after a temporary malfunction of the dissemination system. Third, we impose a maximum threshold for the degree of no-arbitrage violations implied by the option mid-quotes, i.e., the volatility index is n.a. whenever this "non-convexity" threshold is exceeded. Fourth, we use a so-called bounce-back filter to detect short-lived systematic recording errors in the option quotes. It eliminates very few observations, but it guards against errant outliers which induce multiple large artificial jumps in the return series. We loose close to 1.2% of the 15 second observations due to our sequence of filters, with about half of the missing values stemming from the non-systemic staleness requirement and half from the non-convexity condition.[8]

On the other hand, the corresponding 15-second VIX series is obtained from the real-time release by the CBOE. We apply a very mild filter to this series. We remove observations that fall outside the daily high-low range, as reported at the end of trading. This correction reflects errors in the series that have been recognized by the exchange. Second, we also apply the bounce-back filter. We denote the resulting, lightly filtered series VIX$^*$. Figure 1 shows the evolution of the daily VIX closing value over recent years, with the shaded area indicating the sample period exploited in this article.

[6] Thus, we have around 52 (33.5) million OTM put (call) option quotes in the full sample. All trading days are included in the analysis, except for five Holidays with limited trading hours.

[7] For each 15 second cross-section, we exploit an average of 66 (39) strikes for OTM put (call) options for the first maturity, and 69 (43) strikes for OTM put (call) options at the second maturity.

[8] The procedures for implementing these filters are detailed in the Appendix.

Figure 1: **Daily Closing Value for VIX, January 2004 – August 2010**. Our sample period, June 2, 2008 – June 30, 2010, is indicated in green and is shaded.

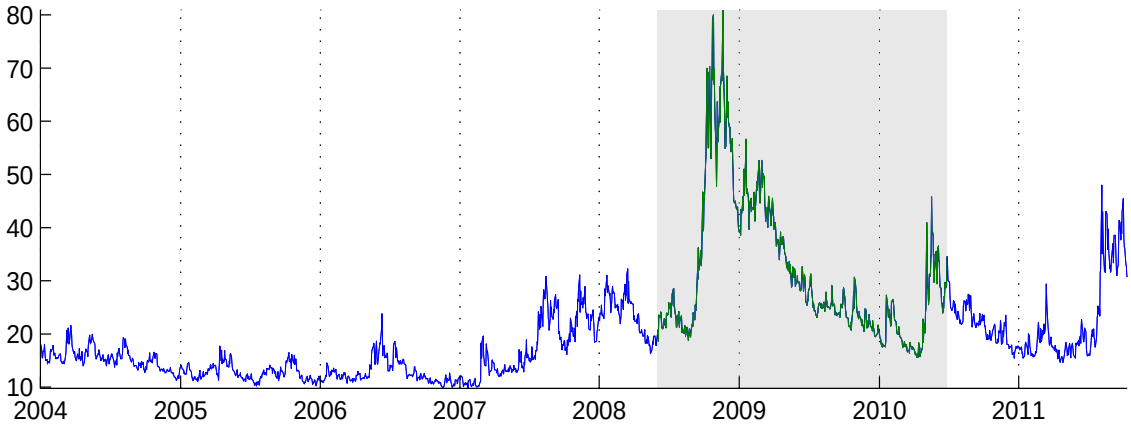


# 4 Practical Computation of the Implied Volatility Indices

## 4.1 The VIX Replication Index, RX1

We first seek to replicate the VIX index from the underlying SPX options using the exact CBOE methodology. The initial step involves computing $\hat{\sigma}^2$. However, the forward price, $F$, and the first strike price below the forward price, $K_f$, in equation (4) are not directly observable. The CBOE determines these variables via put-call parity using the strike price for which the distance between the quoted midpoints of the call and put prices is minimal. We label our version of the VIX index, constructed from SPX option quotes using equations (4)-(5), our "Replication indeX" or RX index. It constitutes our direct equivalent of the CBOE VIX, computed according to official rules.

Our analysis reveals that the forward rate imputation entails a certain lack of robustness, due to the use of only a single put-call option price pair to infer $F$ and $K_f$. Occasional problems, like a temporary gap in the quote updating for a subset of options or a basic recording error, can cause the minimum distance between the call and put prices to, erroneously, appear far from the true ATM strike. In turn, this can generate a large bias in the VIX, and may even render its computation infeasible.

Motivated by this finding, we develop a more robust procedure for determining the implied forward, exploiting all strikes for which the discrepancy between the mid-quote price of the put and call is below \$25. For each of these put-call pairs, we compute the implied forward rate, and then designate the median value as the "robust implied forward." Using a robust statistic prunes the series of major erroneous outliers. On the other hand, the wider set includes less liquid options with larger spreads. Hence, we retain the original forward value, rather than the robust forward, unless they differ by more than 0.5% in relative terms.[9] Thus, our procedure for computing the index is identical to that of the CBOE, except when the implied forward rate deviates substantially from the robust statistic, indicating that the former may contain a sizeable data error. This VIX replication index is denoted RX1.

The left panel of Figure 2 depicts the evolution of VIX$^*$ and RX1 on June 3, 2008. Seemingly, RX1

[9]When the S&P 500 index is around 1000, this threshold is about 5 which corresponds to the gap between individual strike prices. All conclusions in the paper are robust to sensible variation in this threshold.

Figure 2: **Replicating the Real-Time VIX with RX1**. The left panel depicts the VIX* and RX1 indices on June 3, 2008, across the full trading day. The right panel zooms in on a period in the middle of the trading day.

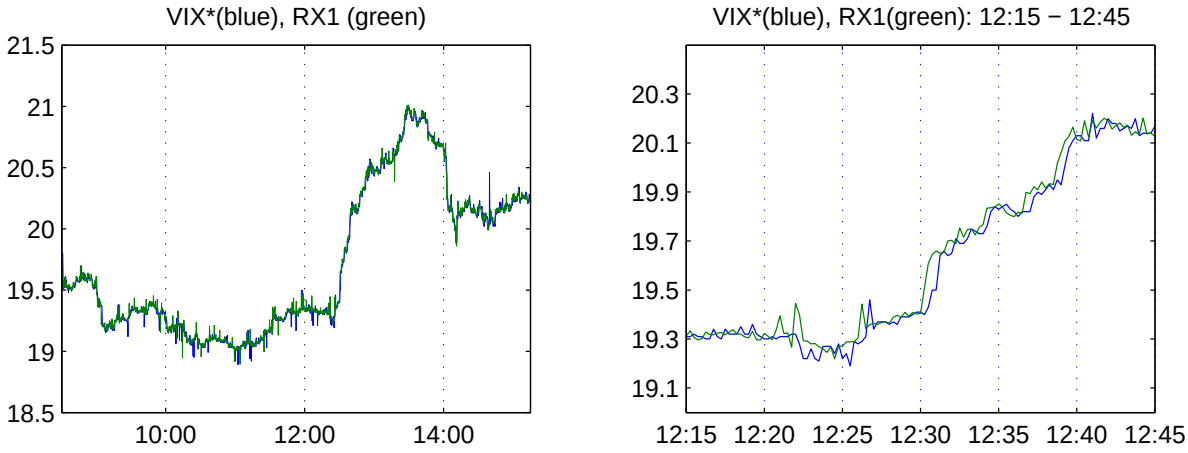


provides near perfect replication, except for minor irregular spikes in the VIX measure not matched by the more smoothly evolving RX1 series. Unfortunately, this is somewhat deceptive. The right panel of Figure 2 zooms in on a shorter time period. It is now evident that VIX lags RX1 by 15-30 seconds. The delay is not unique to this period or trading day, but is observed throughout the sample. It is likely caused by capacity constraints of the CBOE server.[10] Since the reporting delay varies over time in ways that cannot be reconstructed ex post, it is impossible to recreate the record of actual quotes used for the real-time VIX computation. Consequently, it is infeasible to match the high-frequency VIX perfectly, even using the official methodology. Moreover, the relative discrepancies can be non-trivial, even within a fairly calm trading environment. For example, just after 12:30, RX1 rises quickly from 19.4% to above 19.6%. Mechanically, given the delay in the VIX measure, this creates a proportional gap between the series of about 0.75%. Moreover, the deviation becomes more prominent in times of rapidly shifting volatility, when the VIX index is of particular interest as a gauge of market conditions.

These issues may raise concerns regarding the integrity of the real-time VIX series. However, if prices are recorded correctly, albeit with a small lag, the distortion in the general properties of the index will be minor. Our main concerns about the VIX reside elsewhere and run deeper, as explained below.

## 4.2 A Broader VIX Replication Index, RX2

If we ignore the CBOE stopping rule, and exploit all options with positive bid quotes, we obtain a broader strike range, reflecting the *maximum* coverage offered by the prevailing market liquidity. We denote this benchmark index RX2. As for RX1, we use the robust forward price in computing RX2. It should provide an upper bound for the corresponding VIX values. Moreover, whenever there are no positive bid quotes for OTM options beyond the CBOE truncation points, RX1 and RX2 coincide.

It is an empirical matter whether RX1 or RX2 provide the more suitable basis for assessing the shifts in implied volatility over time. As we document later, the answer largely hinges on the relative stability of their strike coverage. We rely on the notion of an *effective range*, or ER, to address this issue. Importantly, this measure controls for the volatility level when gauging the effective coverage of

[10] This explanation arose from conversations with John Hiatt, Director of Research at CBOE.

Figure 3: **RX and VIX Indices, February 16, 2010**. The top panels depict the RX1, RX2 and VIX* volatility indices for February 16, 2010. The bottom panels display the effective strike ranges used by RX1 and RX2 throughout the day, separately for the first (black) and second (blue) maturities.

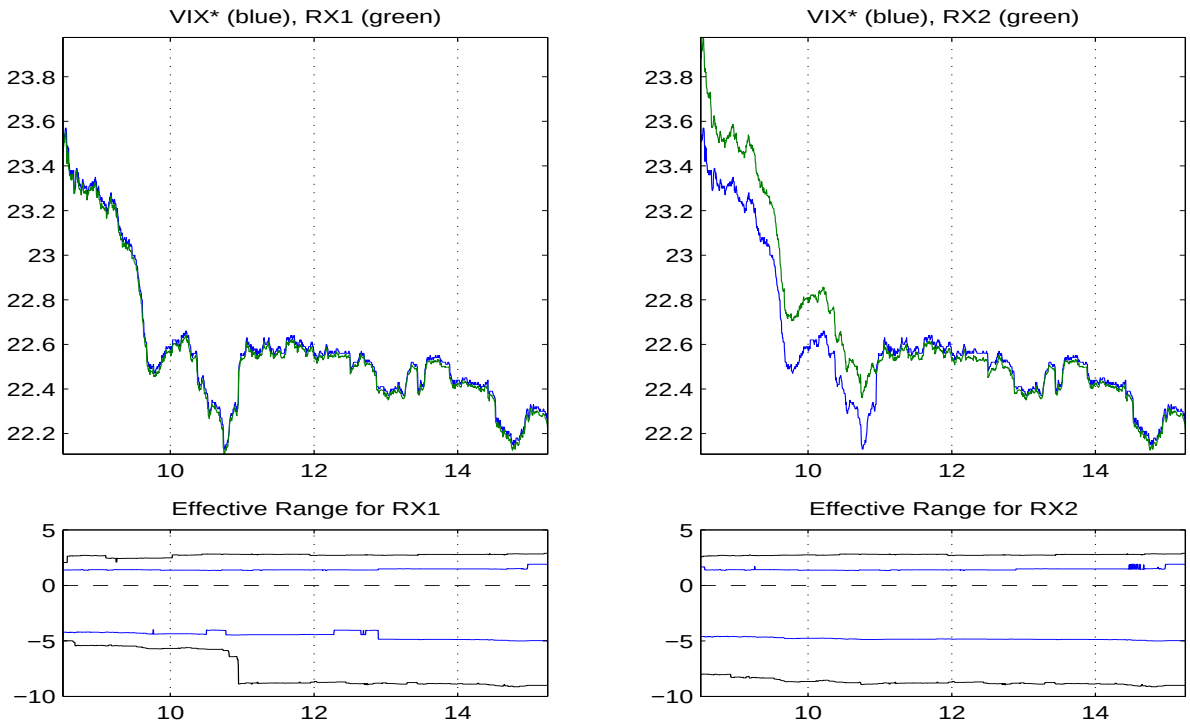


the options in the VIX computation. For the strike range $(K_1, K_n)$, it is defined as,

$$ER = \left[ \frac{\ln(K_1/F)}{\hat{\sigma}_{BS}\sqrt{T}}, \frac{\ln(K_n/F)}{\hat{\sigma}_{BS}\sqrt{T}} \right], \tag{8}$$

where $\hat{\sigma}_{BS}$ is an ATM implied Black-Scholes volatility measure, obtained from a linear interpolation of four Black-Scholes implied volatilities, extracted from options with strike prices just above and below the forward price for the two maturities closest to 30 calendar days.

Figure 3 illustrates why the strike range used to compute the VIX may be of concern. On February 16, 2010, VIX starts out, at 8:30, around 23.5 which is consistent with RX1, but more than two percent below the RX2 value of 24. As the day progresses the gap between VIX and RX2 shrinks to less than half the original size by 10:30, but does not vanish until just prior to 11:00, when both VIX and RX1 "jump" up to RX2. After this point, RX1 and RX2 coincide and replicate the VIX series closely.

At some level, the example represents successful replication of VIX. RX1 closely mimics the real-time evolution of the index. On the other hand, the most notable feature of VIX on this day is the "jump" just prior to 11:00. In a matter of 15 seconds, the index increases by about one percent. However, the RX2 series, computed on the basis of all available options, barely moves at this point – there is no indication of a discontinuity in RX2. The explanation is apparent from the bottom panel of Figure 3. There is a dramatic expansion in the lower part of the strike range (OTM put options at the nearby maturity) used for computing RX1, exactly at the point of the jump. Moreover, there is a drift towards a larger strike range between 8:35 and 10:50 which coincides with the slowly narrowing gap between VIX and RX2 over this period. In contrast, the RX2 measure is computed from a near invariant strike

range throughout the day. As such, RX2 serves as a control, helping to disentangle the effect of changes in the option prices from the shift in the effective strike range. In this case, RX2 clearly provides the more meaningful gauge of the change in (implied) volatility. At least for this trading day, the cut-off rule for the strike range determined by the CBOE produces a suboptimal outcome.

To more systematically assess how well the RX indices replicate the VIX, we compute summary statistics for the "quality of fit" across the full sample. While perfect replication of the high-frequency VIX series is infeasible due to unobserved fluctuations in the dissemination delay at the CBOE, we expect the average discrepancies to be fairly minor. Using a 0.50% (0.25%) relative deviation as the criterion for satisfactory replication, we find that the 15-second VIX$^*$ observations are matched by RX1 about 90% of the time and by RX2 in about 80% of the cases.[11] Of course, since RX2 is not explicitly designed to mimic the VIX, the lower replication rate is not surprising. Overall, the evidence is consistent with a VIX series that oscillates, sometimes abruptly, between the RX2 benchmark and lower values, which may or may not be consistent with RX1. This is what we observed in Figure 3 and it is, indeed, a pattern we recognize across many other days in the sample.

In summary, we generally replicate the VIX index with reasonable precision. Nonetheless, there are troublesome features. Most importantly, the VIX, on occasion, displays large discrete changes solely due to abrupt shifts in the coverage of the underlying options. Moreover, the RX indices may also be subject to such idiosyncratic changes in the amount of active option quotes. The only way to assess the extent of the problem is by explicitly controlling for the effective range used in computing the index.

# 5 Corridor Implied Volatility, or CX, Indices

## 5.1 Defining Corridor Volatility

We turn to the concept of corridor implied volatility, introduced by Carr & Madan (1998) and explored empirically in Andersen & Bondarenko (2007), to assess how variation in the strike range impacts the volatility measures. The point is, a priori, to fix the range of strikes at a level that provides broad coverage but avoids excessive extrapolation of noisy or non-existing quotes for far OTM options. Since it operates with a fairly narrow range, the Corridor Index, or CX, will usually be lower than RX2. However, the invariant strike coverage ensures that the measure is coherent in the time series dimension and alleviates the variation induced by idiosyncratic shifts in the effective range.

The main issue is how to define an economically invariant portion of the strike range which ensures that the associated implied volatility measures are compatible over time. We emphasize a few desirable features. First, the corridor should be linked to observable option prices to ensure transparency and facilitate real-time computation. Second, the forward price constitutes a focal point by determining the ATM strike and thus identifying the set of OTM options to use in the calculation. Thus, it is natural to construct a metric which associates the forward rate with the 50$^{\text{th}}$ percentile of the range. Third, the range should endogenously adjust over time to accommodate variation in the relative importance of strikes in the tail of the distribution. Fourth, it should be model-independent.

To accommodate these features, we adopt the approach of Andersen & Bondarenko (2010). It builds on the fact that option prices reflect **tail moments**. The left and right tail moments of a positive random variable, $x$, with strictly positive density $f(x)$ for all $x > 0$, are given by,

$$LT(K) \;=\; \int_0^K (K-x)\, f(x)\, dx \qquad \text{and} \qquad RT(K) \;=\; \int_K^\infty (x-K)\, f(x)\, dx\,. \tag{9}$$

[11] The result is robust across alternative procedures. Most importantly, we explored the match between VIX$^*$ and 15 or 30 second lagged RX1. The replication frequency rose, but the effect was very marginal.

Figure 4: **Determination of the CX Truncation**. The top left panel displays the normalized OTM option prices at the end of trading on June 16, 2010, for the 30-day maturity. Moneyness is defined as $k = K/F$, while the normalized option prices are $M(k) = Q(T,K)/F$ and $Q(T,K) = \min(C(T,K),P(T,K))$. The top right panel depicts the corresponding Black-Scholes implied volatilities. The bottom panels plot the extracted $R(k)$ function and estimated risk-neutral density. The vertical dashed lines indicate the $1,3,10,50,90,97$ and 99 percentile quotients of $R(k)$.

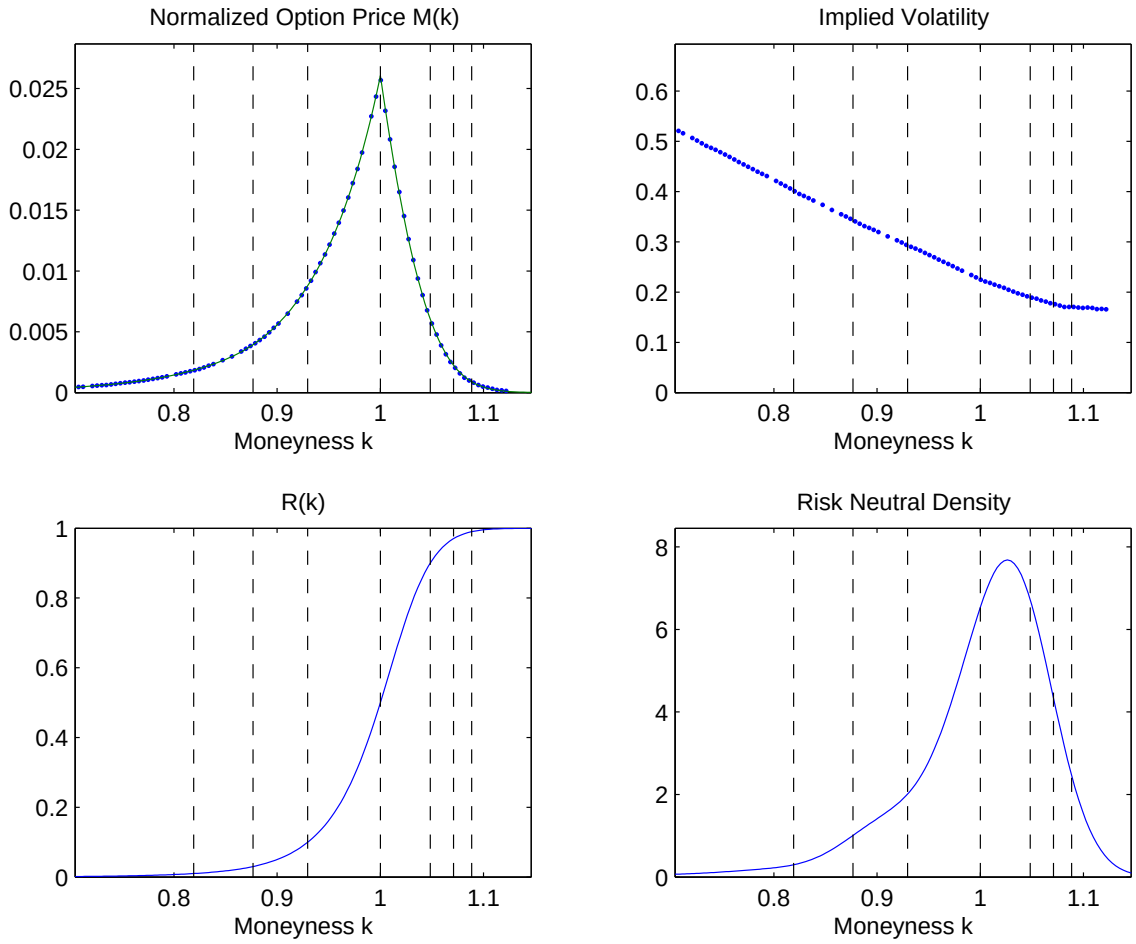


We define the **ratio statistic**, $R(K)$, as an indicator of how far in the tail a given point, $K$, is located within the support of $x$,

$$R(K) \;=\; \frac{LT(K)}{LT(K)+RT(K)}\,. \tag{10}$$

$R(K)$ is akin to a cumulative density function, or $CDF$, as it is increasing on $(0,\infty)$ with $R(0) = 0$ and $R(\infty) = 1$. Moreover, it is centered on the mean of $x$ as $R(\overline{K}) = \frac{1}{2}$, where $\overline{K} = \int_0^\infty x\, f(x)\, dx$. Second, for a given percentile, $q$, in the range of the $R(K)$ function, we define the quotient, $K_q$, as

$$K_q \;=\; R^{-1}(q) \quad \text{for any} \;\; q \in [0,1]\,. \tag{11}$$

Letting $f(x)$ denote the risk-neutral density for the S&P 500 forward price at maturity $T = \frac{1}{12}$, or

one month, and $K$ be the strike price of European style put and call options, we have,

$$R(K) \;=\; \frac{P(K)}{P(K)+C(K)}\,. \tag{12}$$

The ratio statistic is computed solely from put and call prices. Hence, if the corridor is located in the range where option prices may be extracted reliably, $R(K)$ is trivial to compute, and it does not require an estimate of the RND. Moreover, the median, or 50$^{\text{th}}$ percentile, of $R(K)$ corresponds to the expected future value of the equity index, as the forward rate equals the mean of the RND,

$$\overline{K} = F \quad \text{and} \quad K_{0.50} \;=\; R^{-1}(0.50) \;=\; F\,. \tag{13}$$

Consequently, the percentiles of the ratio function conveniently center the strike range on the focal point, $F$, for the computation of model-free implied volatility.

Figure 4 illustrates how option prices across a wide range of strikes are readily deduced from the available quotes. Only towards the tails of the RND are option prices, and thus the distribution function itself, hard to measure accurately. This issue is circumvented by the $R$ function which depends only on the option prices inside the truncation points. It is natural to define the corridor via symmetric percentiles of the $R$ function, so that the truncation points reflect the relative importance of the right and left tails for option pricing in a consistent manner. After exploring the market liquidity across the sample, we settled on a 3% truncation level for the tails. This allows for a broad corridor while ensuring that quotes spanning the width of the corridor are available for almost all the high-frequency intervals. The definition of our corridor variance measure now follows from equation (7),

$$CX \;=\; \frac{2\,e^{rT}}{T}\int_{K_{0.03}}^{K_{0.97}} \frac{Q(T,K)}{K^2}\,dK\,. \tag{14}$$

## 5.2 An Illustrative Trading Day

Figure 5 depicts the evolution of the CX, RX and VIX indices over a volatile day, October 14, 2008, where VIX spans values below 50 and above 60 percent. There is an apparent problem early on: VIX$^*$ is frozen at an inexplicably high level of about 55% from 8:30-8:40, only to later drop to an equally incomprehensibly low value of around 46%. These observations cannot be reconciled with the contemporaneous values for RX1 or RX2 and they are not driven by any unusual activity in the underlying asset prices, depicted in the bottom panels. After additional wild gyrations, the VIX$^*$ values start roughly mimicking the RX1 index in the period 9:40-12:55 and then, after an apparent jump at 12:55, coinciding reasonably well with both RX1 and RX2 until about 14:30. Finally, from this point onwards, RX1 drops below RX2 while VIX$^*$ remains fairly close to RX2 until the market close. The reason for the "jumps" in RX1 at 12:55 and after 14:30 is readily identified from the second panel: the effective strike range for the nearby maturity widens and shrinks in concert with the shifts in the volatility index level. A similar phenomenon occurs for RX2 over 10:05-10:25, when this index indicates a strongly elevated volatility level, while the other volatility indices display no major discontinuity. Again, the event is attributable to an expansion of the associated strike range for the nearby maturity.

In contrast to the other indices, the CX measure evolves continuously, albeit somewhat erratically, throughout the trading day. This suggests that CX provides a more reliable measure of the intertemporal variation in volatility. However, before concluding that the CX index is a superior alternative, we must confirm not only that it avoids some of the high-frequency pitfalls associated with the regular VIX style indices, but also that it captures the broader low frequency volatility movements equally effectively.

Figure 5: **VIX$^*$, RX and CX Indices on October 14, 2008**. The top panels depict the VIX$^*$, RX1, RX2 and CX volatility indices for October 14, 2008. The middle panels display the effective strike ranges used by RX1 and RX2, separately for the first (black) and second (blue) maturities. Finally, the bottom panels indicate the implied forward (left), computed by linear interpolation of the implied forward for the two maturities, and the CME Group S&P 500 futures (right).

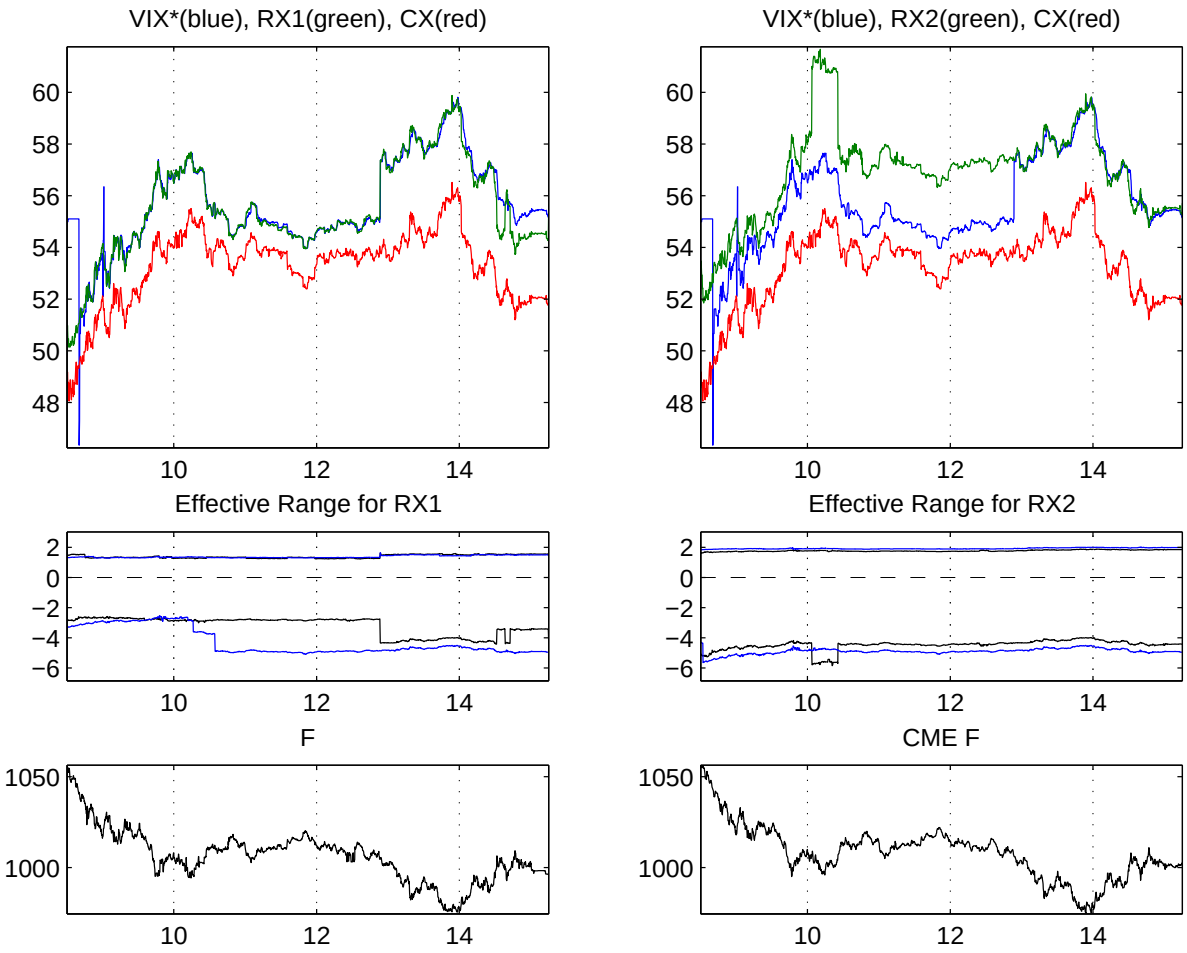


## 5.3 Descriptive Statistics

Option prices are highly sensitive to volatility. Hence, all volatility indices based on a broad cross-section of options should provide similar indicators of the general shifts in volatility. This is manifestly true for the indices explored in this study: all pairwise correlations exceed 0.999. The minimum of 0.9991 is attained for the correlations of CX with VIX, VIX$^*$, and RX1, while that of RX1 and RX2 is 0.9998, and the one between VIX and VIX$^*$ reaches 1.0000. In terms of capturing changes in the (risk-neutral) volatility level at lower frequencies, these measures are effectively identical.

Table 1 provides basic summary statistics for the alternative volatility measures. To alleviate the effect of the small dissemination delay of the VIX measure, the statistics are compiled at the one-minute frequency. Panel A confirms that the series are near indistinguishable in their portrayal of the volatility level. Apart from the fact that the CX index, as a consequence of deliberate truncation, represents a slightly down-scaled version of model-free implied volatility, the statistics are remarkably similar across the board. In particular, the scale-independent skewness and kurtosis measures attain

Table 1: **Summary Statistics for Volatility Indices, 1-min Frequency**. Panel A provides summary statistics for the volatility index levels, quoted in annualized percentage terms. Panel B provides the corresponding statistics for the continuously compounded (log) returns of the volatility indices.

**Panel A: Volatility Index Levels**

| | $RX1$ | $RX2$ | $CX$ | $VIX^*$ | $VIX$ |
|---|---|---|---|---|---|
| Mean | 31.58 | 31.64 | 29.03 | 31.71 | 31.72 |
| Std Dev | 13.28 | 13.33 | 12.31 | 13.42 | 13.42 |
| Skewness | 1.33 | 1.32 | 1.32 | 1.32 | 1.32 |
| Kurtosis | 4.21 | 4.19 | 4.20 | 4.17 | 4.17 |

**Panel B: Volatility Index Returns**

| | $RX1$ | $RX2$ | $CX$ | $VIX^*$ | $VIX$ |
|---|---|---|---|---|---|
| Mean $\times 10^4$ | -0.11 | -0.16 | -0.14 | -0.14 | -0.15 |
| Std Dev $\times 10^4$ | 21.65 | 20.39 | 21.72 | 23.21 | 31.75 |
| Skewness | 0.11 | -0.57 | 0.33 | -0.43 | -0.71 |
| Kurtosis | 54.81 | 90.16 | 21.74 | 378.3 | 3758 |

values around 1.3 and 4.2, respectively, for all the indices. Moreover, the mean and standard deviations are almost perfectly proportional across the measures, buttressing the view that the series, in essence, represent slightly differently scaled versions of the same underlying volatility process.

Of course, our earlier illustrations reveal that the indices can behave very differently on specific trading days. The question is whether this has an appreciable impact on the general high-frequency properties of the series. Panel B of Table 1 provides summary statistics for one-minute percentage changes, or returns, of the volatility measures. We first note the relatively small skewness statistics. There is seemingly little evidence of pronounced asymmetry in the volatility changes, which may be slightly surprising given the commonly adopted assumption that volatility only, or predominantly, jumps upward. We explore this issue extensively later. However, the most striking feature is the huge variation in the kurtosis statistic. While the sample kurtosis for CX is sizeable at around 22, the values for RX1 and RX2 are quite imposing at around 55 and 90, respectively, and truly outsized for the VIX indices. This points towards important differences in the number and size of outliers across the series. The fact that mild filtering reduces the kurtosis of VIX$^*$ dramatically relative to VIX suggests that erroneous data may contribute to the inflated statistics.

Table 2 verifies that the pairwise correlations are dramatically lower for volatility index changes than for the corresponding levels. Instead of near perfect comovement, we now encounter correlations as low as 0.61. The poor coherence between VIX$^*$ and RX1 returns is especially noteworthy, as the latter is designed to mimic the former. Moreover, the VIX returns have the lowest degree of coherence with the "cleaner" CX returns, likely due to the tendency of VIX to randomly oscillate between RX1 and RX2, rendering it an excessively noisy indicator of short term volatility fluctuations. The table highlights the fact that extremely highly correlated series (0.999!) can display widely different high-frequency behavior. For the applications below, this is crucial. Our ability to identify important features of the return dynamics hinges on accurate measures of high-frequency shifts in volatility.

Table 2: **Volatility Index Return Correlations, 1-min Frequency**. The correlation coefficients are computed from all one-minute observations for which (continuously compounded) volatility returns are available for both series.

| | $RX1$ | $RX2$ | $CX$ | $VIX^*$ | $VIX$ |
|---|---|---|---|---|---|
| $RX1$ | 1.000 | | | | |
| $RX2$ | 0.871 | 1.000 | | | |
| $CX$ | 0.856 | 0.895 | 1.000 | | |
| $VIX^*$ | 0.673 | 0.609 | 0.615 | 1.000 | |
| $VIX$ | 0.673 | 0.609 | 0.615 | 1.000 | 1.000 |

# 6 The Dynamics of the Equity-Index Returns

One striking feature of the aggregate U.S. stock market is the tendency for the equity index to occasionally undergo periods of heightened volatility accompanied by predominantly falling prices, while the index at most other times is characterized by lower volatility and a relatively slow, and uneven, ascent in equity prices. The alternation between these "regimes" induces a pronounced negative skewness in medium term equity-index returns. This asymmetry is often attributed to the combination of two factors, namely, first, the strong negative association between innovations in returns and volatility and, second, the tendency for jumps in equity prices to be predominantly negative while concurrent jumps in volatility are positive. Nonetheless, the origin of these effects as well as the most suitable way to capture them via dynamic reduced-form modeling remain the subject of much debate. A major obstacle for resolving these issues is the absence of an observable and reliable spot volatility series that can be used to directly assess the joint high-frequency behavior of returns and volatility.

In this section, we show how the availability of a reliable high-frequency volatility index can help shed new light on these questions. We first provide direct evidence on the frequency of large changes in volatility, finding volatility jumps to be quite common and possessing a near symmetrical size distribution. Next, we investigate the size of the spot correlation between returns and volatility, documenting a striking coherence between estimates across different sampling frequencies. Finally, we illustrate how careful construction of the volatility index allows for a much improved real-time assessment of the market conditions when the trading process is operating under severe stress.

## 6.1 Volatility Jumps

Over the recent decade a number of studies have established the importance of allowing for jumps in volatility for explaining episodes of extreme market returns. A general class of affine jump-diffusion models incorporating both price and volatility jumps is introduced by Duffie, Pan & Singleton (2000). They allow for isolated return or volatility jumps as well as so-called co-jumps, where price and volatility jumps simultaneously. Around the same time, within a high-frequency data setting, Barndorff-Nielsen & Shephard (2001) introduce a volatility process whose innovations consist only of positive jumps. Subsequently, Eraker, Johannes & Polson (2003) show, based purely on an analysis of daily return data, that the evidence for volatility jumps is compelling. This is consistent with the observation that diffusive volatility models featuring jumps in prices alone provide an inadequate fit to option prices, as noted by, e.g., Bakshi, Cao & Chen (1997) and Bates (2000). Finally, Eraker (2004) and Broadie, Chernov & Johannes (2007) explore option pricing for diffusion models that incorporate jumps in both prices and volatility, with the latter finding volatility jumps crucial for a reasonable overall fit. As such,

there is now broad agreement that comprehensive modeling of the return and option dynamics requires flexible specifications of the volatility structure along with jump dynamics for prices and volatilities, see, e.g., Todorov (2010) for more recent refinements along these lines.

By far, the most common way to accommodate jumps in volatility is through specifications that allow for only positive jumps. When co-jumps are considered, this usually involves a strong negative correlation with the return (jump) innovation. The advantage is that one retains near closed-form solutions for asset and derivatives prices as long as the jump intensity and jump size distributions, along with the volatility processes, are given as generalized affine functions of the (volatility) state vector. At the same time, the critical role of these jumps in determining the asset return dynamics inspires new questions. Does volatility only jump upward? How frequent are volatility jumps? Do the distributions for jump frequency and jump size vary over time? Do they depend on the level of volatility? Is the diffusive component of the return dynamics affected by the occurrence of volatility jumps?

The primary reason we have no firm answers to such questions is that we only recently have realized how critical volatility jumps are for the equity return dynamics. This is confounded by the fact that volatility is latent. Drawing inference about jumps in prices, for which we have observable discrete data, is quite difficult. Establishing whether a latent volatility path displays discontinuities from discrete price data alone is, obviously, a more complex undertaking.

The following sections illustrate how our directly observable CX index can be beneficial in establishing new empirical findings of broad relevance for the modeling of volatility jumps. Nonetheless, our empirical analysis is largely exploratory and illustrative. We shall not attempt any comprehensive analysis of the volatility jump dynamics, as this would require an independent full-length study.

### 6.1.1 Identifying Spot Volatility Jumps via Model-Free Implied Volatility

In order to frame our discussion, we impose the weak restriction that volatility follows a separate Ito semimartingale. Hence, return volatility is assumed to evolve according to the following process,

$$d\nu_t \;=\; \xi_{t-}\,dt \;+\; \psi_{t-}\,dB_t \;+\; \kappa_t\,dJ_t\,, \tag{15}$$

where, for any process $y_t$, $y_{t-}$ is the (left) limit of $y_s$ for $s \to t$, $s < t$, $J$ is a jump counting process, so $dJ_t$ takes the value of unity if a jump occurs at time $t$ and zero otherwise, $\kappa_t$ represents the jump size, if a jump occurs at time $t$, and takes values within the set $\mathbb{D}_t \subset \mathbb{R}^0$, which is a, potentially time-varying, subset of the real line excluding zero. Furthermore, $\xi_t$ and $\psi_t$ are the drift and diffusion coefficients. All these processes may be functions of the state vector which may include $\nu_t$, or components thereof, and other economic variables. Finally, the jump compensator (capturing the expected jump contribution) is absorbed into the drift term so the last term in equation (15) is a pure jump process which we, for simplicity, assume to be of finite activity.[12]

The set $\mathbb{D}_t$, indicating the possible jump sizes, warrants a few comments. In principle, we can allow volatility to display negative jumps as long as it remains strictly positive, i.e., we may specify $\mathbb{D}_t = (-\nu_t, \infty) \setminus \{0\}$. However, for many popular models, including the exponentially-affine class, one excludes *any* possibility of negative volatility jumps as it otherwise may be impossible to ensure a positive volatility process. Hence, most studies impose the restriction $\mathbb{D}_t = \mathbb{R}^0_+$. It is ultimately an empirical question whether it is reasonable to constrain the range of feasible jump values in this manner, and it is one of the questions we explore in the following section.

Within the current framework, volatility jumps refer to jumps in spot volatility, $\nu_t$, in equation (1). However, spot volatility is never directly observed or traded and, as argued above, it is hard to

[12] See, e.g., Carr & Wu (2009) for an equivalent representation in a related setting.

construct reliable pathwise realizations of this latent variable from noisy discrete price data. Hence, discontinuities are difficult to ascertain with precision from this type of approach. Instead, we turn to option implied volatility measures as a way to gauge volatility jumps.

All the implied volatility measures we consider are, ideally, functions of variance terms of the form,

$$\mathscr{V}_{t,T}^2 \;=\; \frac{1}{T}\,E^*\left[\int_0^T \nu_u^2\, I_u(B_1,B_2)\,du\right].$$

It follows, in particular, that we can write any of the volatility indices as a smooth function of the state vector which also governs spot volatility. For simplicity, write the index as a function of spot volatility alone, i.e., $\mathscr{V}_{t,T} = H(\nu_t, t)$. Hence, denoting jumps in any variable $y_t$ by $\Delta y_t = y_t - y_{t-}$, we have by Ito's formula for jump-diffusions,

$$\Delta\mathscr{V}_{t,T} \;=\; \Delta H(\nu_t,t) \;=\; H(\nu_t,t) - H(\nu_{t-},t-)\,. \tag{16}$$

This implies that we observe a jump in the volatility index if and only if there is a jump in spot volatility.[13] In other words, we can monitor the high-frequency volatility index for jumps and learn about the jump activity in spot volatility.

On the other hand, there is no simple relation between the jump size in the volatility index and the corresponding jump in spot volatility. In general, the mapping $H$ is a complex function of the state vector, and the relative magnitude of jumps is model-dependent, requiring an explicit characterization of the volatility dynamics. Nonetheless, standard models almost invariably imply monotonicity in the jump size distribution in the sense that large positive jumps in spot volatility are associated with large jumps in the volatility index. Likewise, if negative jumps in spot volatility are accommodated, then any discontinuity in the state vector, inducing a large negative movement in spot volatility, would also manifest itself in a negative jump in the volatility index. In this informal sense, we can gauge the qualitative features of the jump size distribution for spot volatility through the corresponding behavior of the volatility index. At the same time, it is critical that we are sensitive to the potential impact of microstructure noise in the volatility measures so that we avoid a broad mislabeling of the jump activity. Given our earlier evidence, we emphasize the findings obtained from our robust CX index, although we also present selective results from alternative measures. We now turn to this task.

### 6.1.2 Empirical Evidence on Volatility Jumps

This section explores the tendency of alternative volatility indices to display jump-like movements over short time intervals. To alleviate the distortions arising from microstructure noise in the 15-second series, we focus on one-minute changes, or returns. In addition, since the market opening is characterized by occasional failures in establishing adequate liquidity, we compute all statistics in this section from data covering 8:35-15:15 CT, thus excluding the first five minutes of trading each day.

We proceed nonparametrically by defining large moves relative to a robust measure of concurrent volatility. We first obtain a robust estimate of the volatility for each index series across every trading day, accounting both for shifts in volatility across days and the pronounced intraday volatility pattern.[14]

[13] The volatility index may be a function of the full state vector and, in particular, may depend on components of the overall volatility level. Irrespective, the index is functionally related to state variables whose jumps induce simultaneous jumps in spot volatility and the associated volatility index. Moreover, under weak conditions this also applies if the price process contains jumps. See Todorov & Tauchen (2010) for a formal statement of these results in a more general setting.

[14] A description of this procedure is given in the appendix. The qualitative results are robust to a number of alternative ways of estimating intraday volatility.

Exploiting this robust standard deviation measure, we sort the one-minute returns for each volatility measure into a set of mutually exclusive size categories. Table 3 tabulates the findings.

Table 3: **Distribution of Extreme Returns ("Jumps"), 1-min Frequency**. The ranges are measured in multiples of the (robust) standard deviation. The latter is computed as specified in Appendix B.

| | $RX1$ | $RX2$ | $CX$ | $VIX^*$ | $VIX$ |
|---|---|---|---|---|---|
| $(-\infty,-30)$ | 6 | 2 | 0 | 3 | 15 |
| $(-30,-15)$ | 27 | 11 | 1 | 29 | 50 |
| $(-15,-9)$ | 66 | 43 | 21 | 76 | 114 |
| $(-9,-6)$ | 239 | 194 | 128 | 268 | 272 |
| $(-6,-4)$ | 716 | 682 | 613 | 742 | 739 |
| $(4,6)$ | 725 | 655 | 590 | 730 | 731 |
| $(6,9)$ | 247 | 208 | 139 | 251 | 254 |
| $(9,15)$ | 79 | 48 | 10 | 93 | 131 |
| $(15,30)$ | 33 | 12 | 2 | 37 | 56 |
| $(30,\infty)$ | 4 | 2 | 0 | 2 | 11 |

The table reveals a startling discrepancy in the number of large moves across the alternative indices, irrespective of the threshold adopted for identifying "jumps." For example, including moves beyond six standard deviations on the upside and downside, there are approximately 900 outliers in the regular VIX index and 760 for VIX$^*$, while the number is around 700 for RX1 and 520 for RX2. In contrast, it is about 300 for CX. Hence, all indices display, on average, more than one large move every two trading days, and VIX tops the count with well in excess of one per day. Even more telling is the discrepancy in the number of moves beyond 15 standard deviations. The CX measure has three, while the RX2 index has nine times more, and the other indices exceed CX more than twentyfold. This is clearly anomalous: any genuinely large shift in volatility should manifest itself in a significant elevation across a broad range of option prices and should be reflected in any reliable implied volatility index. Thus, these findings instead corroborate the hypothesis that measurement errors, such as idiosyncratic fluctuations in the strike range, seriously inflate the outlier count for many of the indices.[15]

A second intriguing finding is the apparent symmetry of positive and negative "jumps" within each size category of Table 3. A partial explanation is that misclassified jumps naturally revert, as an unusual widening of the strike range, say, often is followed later by a sudden reversal.[16] However, even for measures less prone to such errors, we observe an almost identical number of positive and negative jumps. Table 4 provides formal tests for (unconditionally) symmetry of the CX and VIX$^*$ jump distributions.[17] The results in Panel A confirm that one cannot reject the hypothesis of a symmetric volatility jump distribution for CX. The qualitative findings regarding symmetry of the VIX jumps in Panel B are similar, but this evidence is less compelling as many of these outliers may be spurious.

To further convey the qualitative differences between the extreme moves of our volatility index, CX, versus those associated with the filtered VIX$^*$ index, Table 5 provides the average volatility changes in

[15]The series in Table 3 contain a slightly different number of one-minute observations due to the difference in the filters applied to the series. The qualitative results are unaffected by restricting the sample to the smaller set of one-minute returns that are common to the series. These findings are available upon request.

[16]Our bounce-back filter captures mostly data errors which induce almost instantaneous corrections. Random variation in the width of the strike range typically persists and cannot be filtered without careful analysis of the underlying option data.

[17]The null hypothesis for the binomial test is that a given jump is positive with probability 0.5, so it checks whether the *frequencies* of negative and positive returns are statistically distinguishable. The Kolmogorov-Smirnov test checks whether the *shapes* of the empirical distribution for negative and positive returns are different, apart from the sign reversal.

Table 4: **Binomial and Kolmogorov-Smirnov Tests.** The first two rows tabulate the number of index changes exceeding a specific (signed) threshold. The last two rows provide p-values for the Binomial and Kolmogorov-Smirnov tests for symmetry. The columns indicate the "jump" threshold in multiples of "sigma", computed as indicated in Appendix B.

**Panel A: CX**

| | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
|---|---|---|---|---|---|---|---|---|---|
| Negative | 763 | 326 | 150 | 82 | 38 | 22 | 8 | 5 | 3 |
| Positive | 741 | 317 | 151 | 77 | 36 | 12 | 10 | 7 | 5 |
| $p_{BN}$ | 0.59 | 0.75 | 1.00 | 0.75 | 0.91 | 0.12 | 0.81 | 0.77 | 0.73 |
| $p_{KS}$ | 0.94 | 0.78 | 0.87 | 0.67 | 0.18 | 0.04 | 0.74 | 0.93 | 0.83 |

**Panel B: VIX$^*$**

| | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
|---|---|---|---|---|---|---|---|---|---|
| Negative | 1118 | 601 | 376 | 233 | 162 | 108 | 82 | 71 | 54 |
| Positive | 1113 | 623 | 383 | 248 | 170 | 132 | 101 | 82 | 71 |
| $p_{BN}$ | 0.93 | 0.55 | 0.83 | 0.52 | 0.70 | 0.14 | 0.18 | 0.42 | 0.15 |
| $p_{KS}$ | 0.45 | 0.57 | 0.21 | 0.16 | 0.03 | 0.68 | 0.49 | 0.40 | 0.17 |

both "sigmas" and percentage returns along with the corresponding quantities for the contemporaneous S&P 500 futures returns across size categories. If a large volatility move is spurious, we do not expect the underlying asset to display any unusual patterns at those times, while if there is a genuine large shift in volatility we do expect the equity returns to also be "active." This is, indeed, what the results indicate. The CX "jumps" line up almost perfectly, and negatively, with the equity returns. When we sort the observations in terms of volatility returns from low (large negative moves) to high (large positive moves) we also obtain a monotonic sort of the average equity returns, but of the opposite sign. These effects are greatly muted for VIX$^*$: even if the average VIX$^*$ moves are considerably larger than for CX in the extreme categories, the corresponding absolute equity returns are much lower and the monotonicity across size categories is lost. In short, the CX returns are systematically related to the equity-index returns, while the VIX$^*$ returns display much less coherence with the S&P 500 returns.

Since the existing empirical evidence on volatility jumps is scant, we provide a few additional illustrations. Figure 6 depicts the evolution of CX and the S&P 500 futures on February 12, 2010. We identify a (negative) volatility jump after 11:00am along with a (positive) spike in the equity index at the same time. Although this is the largest volatility move of the day, it barely stands out visually from the figure. In particular, there is no indication of any dramatic market disruption or any data error. Such volatility "jumps" are dispersed across our entire sample, as shown in Figure 7. While there appears to be some mild clustering of these outliers, they are generally distributed quite uniformly.

Our evidence regarding the frequency of "jumps" in return volatility paints a starkly different picture from what is implied by the extant literature. For example, Broadie, Chernov & Johannes (2007) stress the severe finite-sample problem stemming from their finding that jumps are rare – about 1-2 annually – so that even 15-20 years of data is insufficient for precise inference.[18] We identify the same number of jumps over a few days! Of course, this reflects the difficulty in identifying jumps in a latent series from daily data. Discontinuities are, almost tautologically, more readily identified from high-frequency data. With daily data, it is often impossible to detect jumps that are small relative to the daily standard deviation. As a consequence, outliers identified from daily data tend to be much larger than

[18] See their paper for numerous additional studies reaching similar conclusions from daily data.

Table 5: **The Volatility Jump Distribution.** This table categorizes the log returns of the volatility indices CX and VIX* according to size in terms of "sigma," computed as indicated in Appendix B. For each category, the table reports the average jump size, measured in sigma and percent, along with the corresponding averages for the contemporaneous log returns of the S&P 500 futures.

| | *CX* | | *SP* | | *VIX** | | *SP* | |
|---|---|---|---|---|---|---|---|---|
| | $\sigma$ | % | $\sigma$ | % | $\sigma$ | % | $\sigma$ | % |
| $(-\infty,-9)$ | -10.66 | -1.26 | 3.39 | 0.20 | -14.23 | -2.38 | 0.84 | 0.06 |
| $(-9,-6)$ | -7.01 | -0.98 | 2.08 | 0.11 | -7.11 | -1.00 | 1.04 | 0.05 |
| $(-6,0)$ | -0.72 | -0.13 | 0.49 | 0.04 | -0.85 | -0.14 | 0.48 | 0.04 |
| $(0,6)$ | 0.70 | 0.13 | -0.49 | -0.04 | 0.58 | 0.10 | -0.33 | -0.03 |
| $(6,9)$ | 6.99 | 1.05 | -2.79 | -0.17 | 7.07 | 1.10 | -1.08 | -0.07 |
| $(9,\infty)$ | 12.16 | 2.06 | -3.75 | -0.34 | 13.69 | 2.24 | -0.64 | -0.05 |

Figure 6: **The CX Index and the S&P 500 Futures, February 12, 2010** The top panels depict the CX index level and log return across the trading day. The bottom panels display the corresponding series for the S&P 500 futures level and return. For both return panels, the units are expressed in "sigma," computed as indicated in Appendix B.

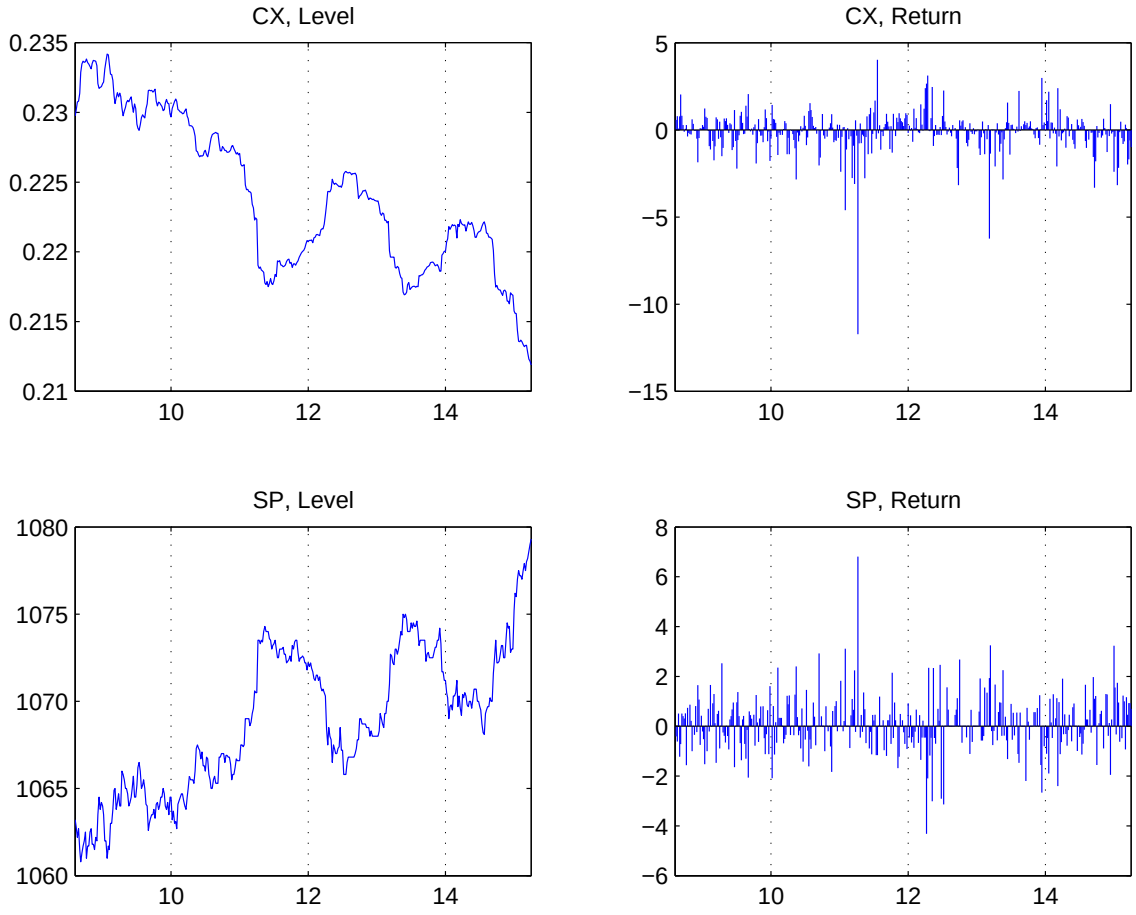

Figure 7: **CX Volatility Jumps** The figure depicts one-minute CX log returns that exceed six "sigmas" in magnitude across the full sample. Sigmas are computed as indicated in Appendix B. The returns exceeding nine sigmas have a circle at the end of their return bar. Positive (jump) return bars are red, while negative return bars are blue.

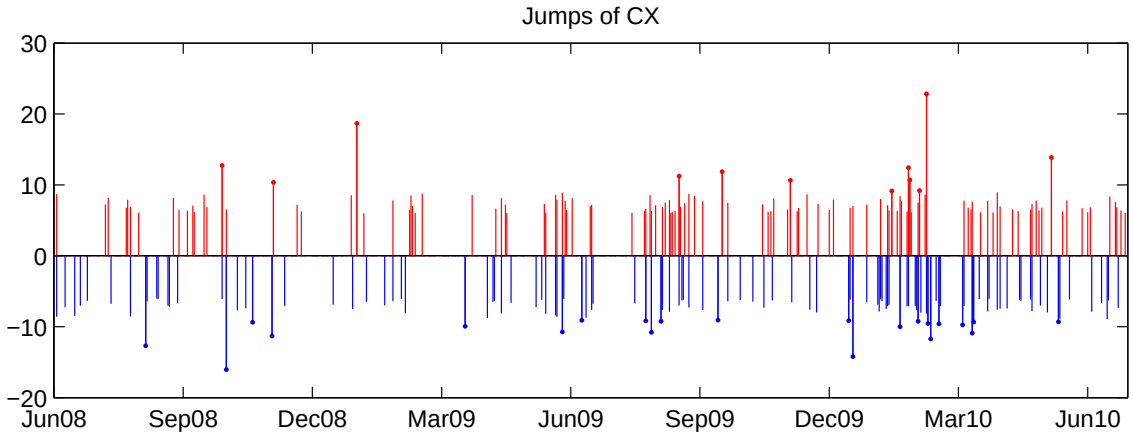


the typical (intraday) jump. As such, the estimated jump intensity is downward biased and the average jump size upward biased. This reasoning suggests that a finer "microscope" is needed to explore the interaction of volatility and price jumps with diffusive return volatility, and how these forces propagate over longer daily and monthly horizons. This study is less ambitious, as we simply provide a few new insights into some critical features of the return dynamics.

In summary, our primary conclusions in this section are, one, that jumps in volatility are more prevalent than recognized in the general asset and derivatives pricing literature. Two, volatility jumps are close to symmetric and, by implication, the evidence for *negative* jumps in volatility is compelling. This is at odds with common parametric specifications of asset price dynamics, including the affine models, which allow only for positive volatility jumps.

## 6.2 The "Leverage Effect"

The asymmetric relation between return and volatility innovations is known as the "leverage effect," following Black (1976) and Christie (1982), who stressed the increased financial leverage, and thus higher return volatility, induced by declining equity prices. The "leverage" label has stuck although, subsequently, this effect has been found to be quantitatively much too weak to rationalize the strong negative association present in the data, see, e.g., Duffee (1995) and Bekaert & Wu (2000). As an alternative, the reverse causality is often invoked: an increase in (non-diversifiable) market volatility may command an increased equity risk premium, inducing a concurrent drop in equity prices. Early contributions along these lines include French, Schwert & Stambaugh (1987) and Campbell & Hentschel (1992). Moreover, recognizing these features, Nelson (1991), Engle & Ng (1993), Glosten, Jagannathan & Runkle (1993) and Harvey & Shephard (1996) incorporate correlations between return and volatility innovations into GARCH and discrete-time stochastic volatility models, typically estimated from daily asset returns. However, investigating the mechanism involved requires access to observations at a much higher frequency as the interaction between price and volatility innovations cannot be disentangled from daily data. Moreover, we would like to develop tools to measure the effect over relatively short samples so that we can establish whether the correlation structure is constant or varies over time in conjunction with market conditions. Consequently, it is advantageous to cast the return

dynamics within a continuous-time setting, enabling analysis at arbitrarily high frequencies.

In light of these considerations a natural definition of the leverage effect is as the limit of the discrete-time return-spot volatility correlation,

$$\mathscr{L}_t \;=\; \lim_{\Delta \to 0} Corr\left(s_t - s_{t-\Delta}\,,\, \nu_t - \nu_{t-\Delta}\right). \tag{17}$$

The leverage effect in continuous time has largely been analyzed in the context of one-factor diffusive volatility models, and in particular within the setting of the so-called Heston model. We now review the relevant issues in this context, before exploring the problems that emerge in multi-factor volatility models and when jumps are present in returns and volatility.

### 6.2.1 The Leverage Effect in Diffusive Volatility Models

A popular continuous-time representation incorporating the leverage effect is the Heston (1993) model. It provides tractable expressions for the discrete-time transition density, so it may serve as a laboratory for analysis of the strength of the return-volatility asymmetry measured at different sampling frequencies, as detailed in, e.g., Bollerslev, Litvinova & Tauchen (2006), Ishida, McAleer & Oya (2011) and Ait-Sahalia, Fan & Li (2012). The model stipulates an affine diffusive evolution for the stochastic volatility factor, which we denote $V = F(\nu)$. $F(\cdot)$ is a smooth, monotone and increasing function of spot volatility, so it accommodates alternative specifications that treat spot volatility ($F(\nu) = \nu$), variance ($F(\nu) = \nu^2$) or log-volatility ($F(\nu) = log\,\nu$) as the basic state variable.[19] The volatility factor is governed by the following stochastic differential equation (SDE),

$$dV_t \;=\; \kappa\,(\theta - V_t)\,dt \;+\; \left(\gamma_0 + \gamma_1\sqrt{V_t}\right)dB_t \quad \text{and} \quad Cov(dW_t\,,dB_t) \;=\; \rho\,dt\,, \tag{18}$$

where, throughout, we assume that the coefficients of the system satisfy regularity conditions ensuring that there is a unique and well-defined solution to the SDE governing the asset return dynamics.

Equation (18) implies that the system has distinct discrete-time transition densities for separate values of the parameter vector and for alternative specifications of the function $F(\cdot)$. Therefore, the correlation of discretely sampled returns and volatilities will depend on both the sampling frequency and the functional form. This induces potential ambiguity into the measurement of leverage. However, we circumvent any such issues by invoking the definition in equation (17), as demonstrated below.

The invariance of the leverage coefficient across alternative specifications is a critical feature that guides our efforts towards a robust estimation procedure. We first document that this property applies for the entire class of one-factor stochastic volatility diffusions, and not only for the Heston model. Towards this end, we recall that the return dynamics, as indicated in equation (1), is given by,

$$ds_t \;=\; \mu_t\,dt \;+\; \nu_t\,dW_t\,,$$

where $\nu_t$ captures the stochastically evolving volatility process. As above, we specify the dynamics for a function of volatility, $V_t \;=\; F\left(\nu_t\right)$. It follows a general one-factor diffusion,

$$dV_t \;=\; \xi\left(V_t\right)dt \;+\; \psi\left(V_t\right)dB_t \quad \text{and} \quad Cov\left(dW_t\,,dB_t\right) = \rho_t\,dt\,. \tag{19}$$

The drift and diffusive coefficients may be arbitrary smooth functions of the underlying spot volatility. Note that the Heston dynamics is obtained, as a special case, by specifying the drift and (squared) dif-

[19] For example, Bandi & Reno (2012) explore the leverage effect using a log-volatility specification, while Wang & Mykland (2012) analyze the same general representation that we employ here.

fusive coefficients as affine functions of the volatility state and imposing a constant correlation between the return and volatility innovations, $\rho_t = \rho$. Moreover, in the following, we frequently suppress the explicit dependence on the state vector and write, e.g., $\psi_t = \psi(V_t)$.[20]

The leverage coefficient for the general one-factor diffusive volatility model is now readily computed via Ito's Lemma. We allow the volatility to be represented by any smooth and monotone increasing function of the underlying volatility factor, $H(V_t)$, with derivative $H'_t = H'_t(V_t)$,[21]

$$\mathscr{L}_t = \frac{Cov(ds_t, dH(V_t))}{\sqrt{Var(ds_t) \cdot Var(dH(V_t))}} = \frac{\nu_t H'_t \psi_t \rho_t}{\sqrt{\nu_t^2 \cdot (H'_t)^2 \psi_t^2}} = \rho_t . \tag{20}$$

The one-factor assumption is restrictive. Empirical studies often conclude that diffusive volatility models should contain at least two stochastic volatility factors – one to accommodate relatively abrupt but quickly mean reverting fluctuations in volatility, and one to account for the pronounced persistence observed in long-run volatility. If additional factors are included, we obtain a model approximating long-memory behavior in volatility along the lines of Corsi (2009), see, e.g., Andersen & Bollerslev (1997). It is straightforward, albeit notationally more involved, to repeat the analysis for a specification including multiple independent volatility factors. One may then show that the leverage coefficient continues to be uniquely identified from the spot return-volatility correlation. Moreover, even if each volatility factor displays an invariant degree of correlation with the return process, the aggregate return-volatility correlation will generally fluctuate over time, as the volatility factors evolve.[22] The only exception is the case where all volatility factors have the identical correlation coefficient vis-a-vis the returns. The upshot is that we can estimate the overall spot return-volatility correlation but we cannot identify the underlying correlations for individual volatility factors in a model-free fashion.

In conclusion, the leverage is uniquely defined by the spot correlation between the return and volatility innovations, irrespective of the functional form for the volatility process or any other parameter governing the system. In particular, while the discrete-time transition densities for the Heston model depend on all features of the system, the leverage coefficient is invariant and equal to $\rho$. However, in general, for realistic specification we expect the leverage coefficient to display some fluctuations over time, both because the effect intrinsically may depend on the general economic environment and because there may be multiple volatility factors, each with their own degree of correlation with the underlying returns.[23] These observations reinforce the point, consistent with definition (17), that we, ideally, should estimate leverage via the return-volatility correlation in high-frequency data. Otherwise, the estimator is subject to biases depending, in potentially complex and unknown ways, on the specific form of the $F(\cdot)$ function and the exact (lower) sampling frequency. Moreover, we can only capture time-varying leverage effects to the extent that we can estimate the correlation over short time periods.

### 6.2.2 The Leverage Effect with Price and Volatility Jumps

Another important consideration is the potential impact of major price and/or volatility jumps on the leverage effect. It depends to a large extent on the interaction between the price and volatility jumps. In

[20] In this shorthand, the notation for the drift and diffusive coefficients is similar to that in equation (15). Since the role of the coefficients is broadly similar, this is useful, but the coefficients should not be treated as identical across those equations.

[21] Again, this allows us to define the leverage effect corresponding to the spot volatility, variance or log-variance, or even a more general volatility index that is a function of spot volatility.

[22] This result is discussed in detail by Bollerslev, Litvinova & Tauchen (2006). The logic is similar to the case when a jumps are included in the system. We analyze that scenario in the following section.

[23] See Veraart & Veraart (2012) for a novel approach incorporating independent and persistent fluctuations in leverage.

particular, as we demonstrate below, an isolated jump in price or volatility reduces leverage. In contrast, the impact of a simultaneous jump in price and volatility hinges on the dependence between the two jump distributions. The latter case is of significant interest as there is evidence for the presence of such "co-jumps" in equity-index prices and volatilities, see, e.g., Jacod & Todorov (2010). Consequently, we need to be specific about the properties of the jump processes and allow for alternative scenarios.

For illustrative purposes, we posit a simple one-factor volatility model with both isolated price and volatility jumps as well as co-jumps of finite activity, denoted $J_t^s$, $J_t^v$, and $J_t$, respectively.[24] To avoid excessive notation and cleanly identify the main features, we assume all jump size distributions are invariant and independent of the volatility state, with mean jump size zero and constant variance. We allow for time-variation in the jump intensities, and we assume constant return-volatility innovation correlations for the diffusive and jump terms. This leads to the following simplified system,

$$\begin{aligned} ds_t &= \mu_s(v_t)\,dt + v_t\,dW_t + \kappa_s\,dJ_t^s + \kappa_J\,dJ_t\,, \\ dv_t &= \mu_v(v_t)\,dt + \psi_t\,dB_t + \delta_v\,dJ_t^v + \delta_J\,dJ_t\,. \end{aligned}$$

The (isolated) log-price and volatility jump sizes have constant variances of $\sigma_\kappa^2$ and $\sigma_\delta^2$, while the sizes of the co-jumps have constant variances of $\sigma_{J\kappa}^2$ and $\sigma_{J\delta}^2$, respectively. The (finite) jump intensities are, $\lambda_{s,t} = \lambda_s(v_t)$, $\lambda_{v,t} = \lambda_v(v_t)$, and $\lambda_{J,t} = \lambda_J(v_t)$. Finally, we denote the correlation between the return and volatility jumps by $\rho_J$, and we label the diffusive correlation, $Corr(dB_t, dW_t) = \rho_C$.

Direct computations now yield,

$$\mathscr{L}_t = \frac{v_t \cdot \psi_t \cdot \rho_C + \lambda_{J,t} \cdot \sigma_{J\kappa} \cdot \sigma_{J\delta} \cdot \rho_J}{\sqrt{\left(v_t^2 + \lambda_{s,t} \cdot \sigma_\kappa^2 + \lambda_{J,t} \cdot \sigma_{J\kappa}^2\right) \cdot \left(\psi_t^2 + \lambda_{v,t} \cdot \sigma_\delta^2 + \lambda_{J,t} \cdot \sigma_{J\delta}^2\right)}}\,. \tag{21}$$

Equation (21) shows that the jump scenario generally implies a time-varying leverage effect, as the return-volatility correlation is a linear combination of the continuous and jump leverage coefficients with weights depending, in potentially complex ways, on the level of volatility.

A number of cases of interest may be identified. If there are no jumps, i.e., $\lambda_{s,t} = \lambda_{v,t} = \lambda_{J,t} = 0$, we obtain $\mathscr{L}_t = \rho_C$, as noted previously. If there are no co-jumps, i.e., $\lambda_{J,t} = 0$, the presence of isolated jumps in either price or volatility drives the leverage coefficient towards zero, as the numerator is unchanged from the pure diffusive setting, but the denominator increases. If only co-jumps are present, i.e., $\lambda_{s,t} = \lambda_{v,t} = 0$, the impact of jumps hinges primarily on the relative size of the diffusive and jump correlations, $\rho_C$ versus $\rho_J$.

Importantly, the leverage coefficient in equation (21) refers to an expected relationship, summarized by the correlation, which reflects the jump intensity and the expected squared jump sizes conditional on jumps occurring. For a given (short) trading period, jumps may or may not be present. This suggests a way to gauge the impact of jumps. If we sort the high-frequency observations into separate categories reflecting the size of the (absolute) change in volatility or price, we can test whether the realized correlations during periods containing large volatility or price changes (jumps) are different from those when no such "jumps" are present. If isolated jumps are frequent, the "jump sample" should display significantly less leverage than the remaining sample. In contrast, if co-jumps are frequent, the "realized correlation" for the sample with large jumps will be indicative of the size of $\rho_J$.

[24] A similar setting is explored by Bandi & Reno (2012) although they define leverage differently and do not exploit high-frequency volatility index observations.

### 6.2.3 Model-Free Identification of the Spot Return-Volatility Correlation

The previous sections suggest the use of high-frequency data may be critical for estimation of the leverage coefficient. Nonetheless, there are reasons to be wary of reliance on intraday return data alone. For example, Ait-Sahalia, Fan & Li (2012) document significant biases due to the fact that (spot) volatility is latent and must be estimated. First, individual returns are obtained by sampling over discrete intervals (discretization bias). Second, distortions from market microstructure noise imply that such intervals must be of non-trivial size. Third, the squared returns are smoothed in a small neighborhood to produce a reasonably accurate local volatility estimate (using integrated rather than spot variance), resulting in even longer effective sampling intervals. Fourth, even after the smoothing, the estimation error for the local (spot) variance remains non-negligible. Instead, Ait-Sahalia, Fan & Li (2012) propose estimating the correlations over various multi-day horizons and using a simple regression framework to infer the underlying spot correlation coefficient from the specific theoretical relationship among these values for the Heston model. They confirm that their procedure avoids the tendency for the leverage coefficient to be biased strongly upward towards zero, as typically happen when estimating directly from high-frequency return-volatility correlations.

In spite of encouraging empirical results, the approach of Ait-Sahalia, Fan & Li (2012) is restrictive along important dimensions. Firstly, the one-factor Heston model does not provide a satisfactory fit to the return dynamics, so it is awkward to impose specific features of this model in inferring the leverage coefficient. Secondly, the procedure rules out time variation in the leverage coefficient and, by construction, cannot shed light on the issue. Thirdly, more suitable volatility models do, indeed, point towards the importance of allowing for a time-varying leverage effect.

As noted, the main complicating feature is that spot volatility is not observed. This points to a possible solution: exploiting our robust model-free implied volatility measure, CX. This index allows us to circumvent many of the obstacles. First, the index is based on option prices, so there is no need for estimation of spot volatility. Second, we have high-frequency observations available so there is no need for any time aggregation. Third, the index is less noisy than the regular high-frequency VIX series. Fourth, we can monitor the effects of microstructure noise and select a sampling frequency that is reasonably immune to such distortions. In combination, these features allow us not only to alleviate the estimation, aggregation and discretization biases, but also to estimate the leverage coefficient at a daily – or even shorter – horizon to address the issue of potential time variation.

An oft-cited concern is that the CX index represents a 30-day *risk-neutral* volatility measure, so it differs from the spot volatility under the physical measure. However, as discussed previously, the equivalence of the physical and risk-neutral measures ensures that we can identify the spot correlation coefficient under weak assumptions. Specifically, in the one-factor diffusive volatility case, the expected future risk-neutral return variation, very generally, is a smooth function of a volatility-related state variable. Since spot volatility is identical across equivalent measures, we have $CX_t = H(V_t)$ for some smooth mapping, $H(\cdot)$. Hence, the result in equation (20) applies, and the relevant correlation coefficient may be estimated directly via the high-frequency return-CX correlation. Likewise, if there are multiple volatility factors, the components of volatility will be smooth functions of the state vector and the basic result goes through. Of course, the same logic applies to the regular high-frequency VIX measure, although reliance on this series may be problematic given the degree of noise in that index.

### 6.2.4 Empirical Evidence on the Return-Volatility Asymmetry

The pronounced negative correlation between daily S&P 500 and VIX returns is well established and serves as an argument for diversifying long equity positions with an exposure to the volatility index.

Figure 8: **Signature Plot for the Volatility Index–S&P 500 Return Correlation**. The left (right) panel depicts the volatility index correlation with the forward (futures) return across frequencies; CX is blue, VIX* is green.

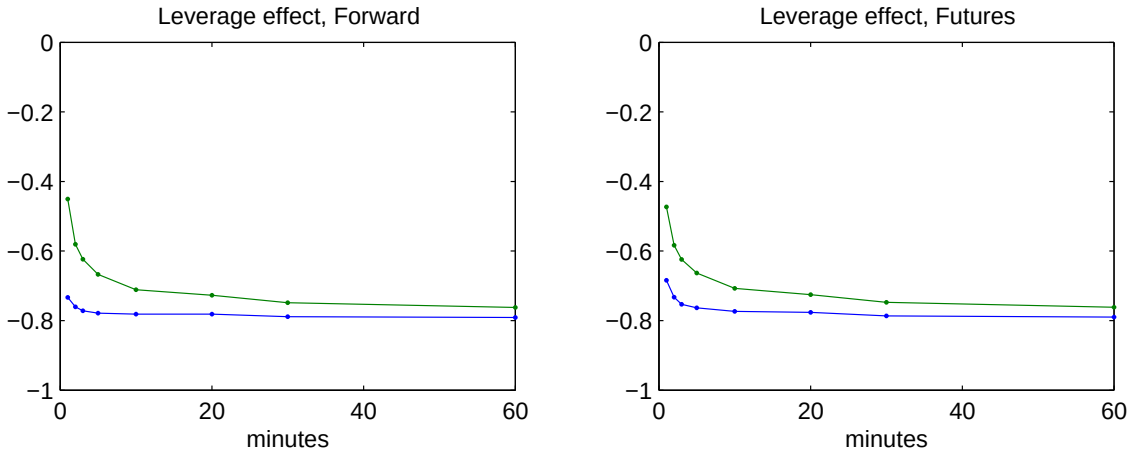


The CBOE web-site provides year-by-year estimates for the sample correlation of the two series in the range of –0.75 to –0.85 at the daily level for 2004–2009. The –0.75 estimate refers to the year 2009 which constitutes about half of our sample. The origin of such large negative correlations has been much debated. For example, it is unclear whether it arises from a corresponding correlation at the very highest frequencies or, alternatively, from a feedback mechanism where an elevation of volatility raises the required rate of return on the equity index, inducing a drop in stock prices. Moreover, these features may interact, inducing renewed volatility spikes and price declines. Such effects are hard to ascertain from daily data, as the aggregation obscures the underlying real-time dynamics.

An important initial robustness check is to establish whether the extreme biases in the return-volatility correlation at high frequencies, documented by Ait-Sahalia, Fan & Li (2012), are present in our context. Figure 8 provides a so-called signature plot, displaying the estimated (unconditional) return-volatility correlations for the full sample against the underlying intraday sampling frequency. The estimates for the CX–S&P 500 return correlation are consistent across the frequencies, with only the results based on the one- and two-minute returns displaying a hint of any bias. This supports the notion that a carefully calibrated option-implied volatility index allows us to sidestep the problems associated with the latency of the return volatility. Given the efficiency gain stemming from using a larger number of observations, especially for estimating correlations over short horizons, we focus on results based on the one-minute returns below. Not surprisingly, the correlations based on the VIX* index are less stable and never attain the low value associated with the CX series. Nonetheless, even for this series, only minor biases appear present for sampling frequencies below 20 minutes. Finally, it seems preferable to exploit the implied forward rather than the futures returns at the highest sampling frequencies. This is likely due to imperfect data synchronization between the two markets.

Next, to gauge whether the presence of artificial jumps distorts the results, we split the observations into three sets, reflecting the size of the volatility returns relative to the robust estimate of the standard deviation, namely (0,6) (RM, or "Regular Move"), (6,9) (SJ, or "Small Jump") and (9,∞) (LJ, or "Large Jump"), thus exploiting the classifications from Table 5. The rows of Table 6 report S&P 500 volatility-return correlation estimates for alternative volatility indices across these three size categories along with the overall correlation estimate.

Table 6 conveys a clear message. For "small" volatility returns, the correlations with the equity

Table 6: **Correlation of Volatility Indices with S&P 500 Returns.** The S&P 500 returns are based on the implied forward prices. The correlations are computed from one-minute log returns.

| | $RX1$ | $RX2$ | $CX$ | $VIX^*$ | $VIX$ |
|---|---|---|---|---|---|
| $(0,6)$ | -0.68 | -0.70 | -0.73 | -0.48 | -0.48 |
| $(6,9)$ | -0.57 | -0.57 | -0.74 | -0.49 | -0.49 |
| $(9,\infty)$ | -0.38 | -0.45 | -0.76 | -0.28 | -0.11 |
| $(0,\infty)$ | -0.64 | -0.67 | -0.73 | -0.45 | -0.33 |

returns are uniformly strongly negative, coming in at around –0.70 for the CX and RX series and around –0.50 for the VIX indices. For the small jumps, in the $(6,9)$ category, the correlation is equally negative for CX, but falls noticeably for the RX indices. Finally, for the larger jumps, only the CX correlation remains invariant while it drops off dramatically for the other series. In summary, the CX index stands out by displaying a consistent degree of correlation with the equity returns.[25]

Obviously, the correlations in Table 6 are estimated with some error. Table 7 provides bounds on the estimates for the CX and VIX$^*$ indices, reflecting the uncertainty associated with pure sampling variability.[26] It is evident that the correlations for the full sample as well as the regular moves are estimated with high precision. Since the jumps are much less common, the associated correlation estimates are considerably more imprecise. Hence, even if the strength of the return-volatility asymmetry appears close to identical across the size categories, this cannot be verified without access to an even larger sample. In addition, we caution that the one-minute return correlation estimates may be slightly biased towards zero, as suggested by the pattern observed in Figure 8. Thus, our –0.73 point estimate likely constitutes a conservative upper bound for the actual (unconditional) correlation coefficient. The true (average) correlation coefficient is probably slightly more negative.

We note that the corresponding results for VIX$^*$ in Table 7 are irreconcilable with the above findings, based on our CX series. The VIX correlations point towards much less asymmetry in the average return-volatility relation along with a sharply diminished leverage coefficient when volatility undergoes abrupt shifts. Of course, we have already shown that the high-frequency volatility measure released by the CBOE is plagued by an excessive degree of idiosyncratic noise and contains "artificial jumps." The present findings suggest that a direct application of the high-frequency VIX series for assessment of the leverage effect will produce badly distorted inference. Even sensible ex-post filtering, as embedded in the VIX$^*$ series, is insufficient to remedy the problem.

The highly negative value of our point estimate for the leverage coefficient signifies another substantial departure from the extant literature. For example, in work allowing for time-varying leverage, Bandi & Reno (2012) find the coefficient mostly ranging from –0.22 to –0.35, and only reaching –0.45 for exceptionally high levels of realized volatility. Ishida, McAleer & Oya (2011) exploits high-frequency VIX data directly and obtain alternative GMM based estimates spanning –0.51 to –0.57. Finally, Ait-Sahalia, Fan & Li (2012) end up with a point estimate of –0.676 after correcting for numerous effects that tend to bias the coefficient towards zero. Consequently, our approach produces a more dramatic estimate for the strength of the contemporaneous equity-index return-volatility asymmetry than obtained in recent studies dedicated exclusively to inference regarding the leverage coefficient. It is also striking that our estimate, obtained from the high-frequency CX series, is in line with the

[25]The findings based on the correlation with the futures returns are qualitatively similar, but the estimates are more dispersed. These results are available upon request.

[26]The confidence bands are constructed as prescribed by Fisher (1921).

Table 7: **S&P 500-Volatility Return Correlation Bounds.** The table provides point estimates as well as upper and lower bounds of the 95% confidence bands for the CX and VIX* log return correlations with the (implied forward) S&P 500 log return.

| | *CX* | | | *VIX** | | |
|---|---|---|---|---|---|---|
| | Lower bound | Point estimate | Upper bound | Lower bound | Point estimate | Upper bound |
| $(0,6)$ | -0.737 | -0.735 | -0.733 | -0.488 | -0.484 | -0.481 |
| $(6,9)$ | -0.793 | -0.744 | -0.685 | -0.555 | -0.491 | -0.423 |
| $(9,\infty)$ | -0.873 | -0.759 | -0.567 | -0.395 | -0.277 | -0.150 |
| $(0,\infty)$ | -0.735 | -0.733 | -0.731 | -0.454 | -0.450 | -0.447 |

value obtained from daily VIX and equity returns over the corresponding sample period – even if this correspondence is not directly implied by the underlying theory.

We conclude the section by briefly exploring whether there is any indication of time variation in the leverage effect. For that purpose the leverage coefficient is estimated for every trading day in our sample using one-minute futures returns from the CME group and our one-minute CX returns derived from the CBOE options data. The reliance on data from separate exchanges avoids the potential for a mechanical distortion in the correlation estimates due to errors simultaneously impacting the option and equity-index prices. These estimates are then averaged using a 21-day moving centered window, so the series provides a sequence of daily overlapping measures of the monthly leverage coefficient. The smoothing helps alleviate noise and more clearly bring out trend movements. A simple indicator of market conditions over the sample is provided by our implied volatility index, CX. Figure 9 overlays the two series. First, we note that monthly leverage evolves within a fairly narrow range of about $-0.68$ to $-0.91$. Thus, even our highest monthly estimate is below the one obtained by Ait-Sahalia, Fan & Li (2012), and it is well below other typical estimates in the literature. Second, there appear to be cyclical patterns in the series. The leverage increases as the CX index declines. This is most evident in trend-like behavior of the leverage coefficient over particular sub-periods: it is rising over June-August 2008, declining over September-December 2008, when it hits the low point of the sample, increasing and then relatively flat over January-July 2009, but with notable dips during the smaller spikes in CX during February and early March, and then generally quite elevated until June 2010. Finally, there is a distinct drop around the renewed market turmoil in the Summer of 2010.

Figure 9 suggests that the strength of the leverage effect is linked to the level of implied volatility. The correlation between the series is $-0.588$. Although the sample is limited, preventing us from drawing firm conclusions, the finding is intriguing. Periods with high implied volatility have been found to involve an unusually large variance risk premium and an elevated equity risk premium. Hence, the increasing (absolute) leverage coefficient during such scenarios suggests that the effect is associated with risk pricing.[27] During "bad" economic times, when market conditions generally are turbulent, the negative association between equity-index prices and volatility seems to grow even stronger. Since we are exploiting one-minute observations, we cannot conclude that the effect is truly instantaneous, but these return-volatility interactions are both swift and pronounced.

While the association between the leverage effect and concurrent market conditions may be natural, there is no a priori reason to believe the relation is linear. In particular, the implied volatility index is strongly right-skewed while the logarithmic CX series is more homogeneous. Figure 10 provides

[27] See, e.g., Nagel (2012) for a broader discussion of the evidence relating implied volatility and risk premiums.

Figure 9: **CX and Leverage.** CX is blue, the leverage coefficient is green. The leverage effect is computed with respect to the S&P 500 futures and it is averaged over a moving centered window of 21 days.

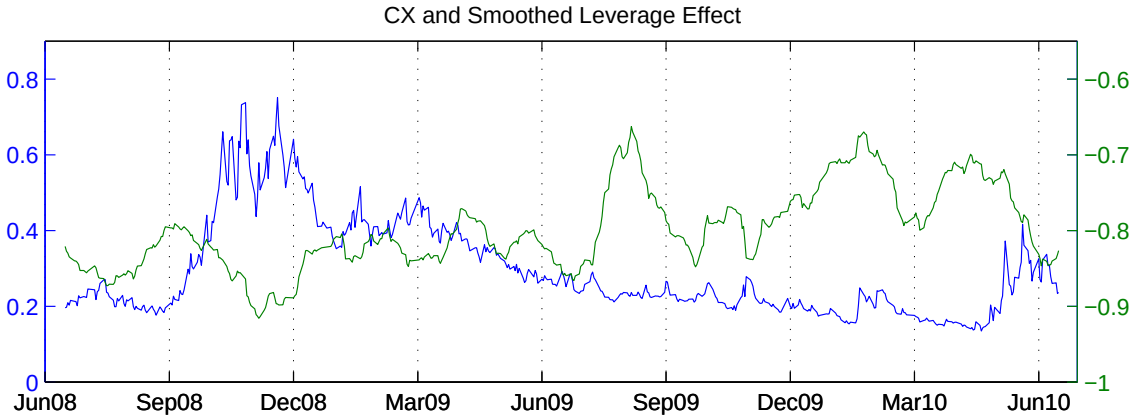


Figure 10: **Scatter Plot of Leverage versus CX and log CX.** The leverage coefficient is computed with respect to the S&P 500 futures and it is averaged over a moving centered window of 21 days.

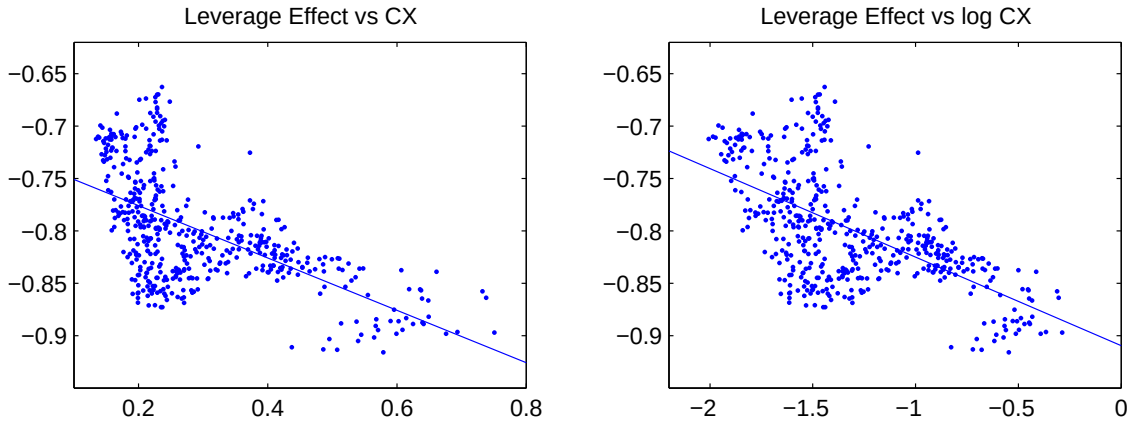


scatter plots of the relation between the monthly leverage measures and both the CX and the log CX index levels. While the log transformation does alleviate the clustering of volatility across the range of the x-axis, there is no striking difference in the overall quality of fit across the two panels, so the evidence for a pronounced negative relation appears robust.

We emphasize that the ability to measure the leverage coefficient relatively accurately at a high frequency is key to our analysis. This is, seemingly, infeasible to accomplish using a methodology relying exclusively on equity-index returns.[28] This opens up avenues for future empirical and theoretical work exploring the equity return-volatility asymmetry in more depth and detail.

[28] Bandi & Reno (2012) also find evidence for an increasingly strong leverage effect during high volatility periods, but their average estimates are much less pronounced than implied by our results. However, for return-volatility co-jumps, Bandi & Reno (2011) report estimates at the other extreme, namely very close to $-1$.

### 6.3 Volatility Indices During Periods of Market Stress: The Flash Crash

On May 6, 2010, just after 13:30 CT, the S&P 500 index dived by more than five percent in a matter of minutes, only to rapidly recover thereafter. Most other U.S. securities markets underwent corresponding whipsawing trading patterns, pointing towards a breakdown in cross-market liquidity. This decline and rebound was unprecedented in speed and scope. The event is now known as the "flash crash." In the CFTC-SEC Report, exploring the market dynamics of this episode, the intraday VIX measure features prominently in the account of the escalating uncertainty gripping the market.

The top left panel of Figure 11 displays the VIX and E-mini S&P 500 futures price for that trading day.[29] Given the potential problems with VIX during turbulent market conditions, we also plot the S&P 500 series against a corridor volatility index, $CX^*$, which equals our CX measure, except that it has been scaled to match the general level of the VIX index. The middle panels depict the VIX versus RX1 and CX, respectively. Finally, the bottom panels show the effective strike range used in the computation of RX1 and the relative discrepancy between VIX and CX across the day.

A number of features in Figure 11 are striking. First, during the early part of the trading day, the VIX and $CX^*$ indices are near indistinguishable and both hit a local peak at the trough of the equity-index around 13:46. Second, both decline in erratic fashion over the following 20-25 minutes. Third, the indices then diverge sharply over the next hour, until after 15:00. The VIX spikes up again and attains its maximum around 14:30, while CX is more than ten percent lower and well below its high of the day. Moreover, the path of the VIX is now characterized by abrupt jumps, while CX evolves more smoothly, almost as a mirror image of the S&P 500 index. Hence, the two series portray vastly different trading environments following the crash. The real-time VIX index conveys a picture of sustained, or even elevated, uncertainty for a prolonged period. The corridor index instead settles down reasonably well around 14:00, and subsequently never reaches a level comparable to the peak at 13:46.

For a real-time observer, the successive VIX spikes in excess of 10% around 14:12 and 14:30 would be rather unsettling. However, if the CX series is reliable, these volatility eruptions were fictitious. The stark divergence between the measures is due to features that are, by now, familiar. The left middle panel shows that RX1 mimics VIX very well throughout the trading day. In other words, we can rationalize all major shifts in the VIX from the underlying option quotes and the CBOE truncation rule. At the same time, the bottom left panel reveals dramatic instability in the effective strike range used to compute the VIX. By construction, our CX measure is robust to such problems. Hence, the relative value of the VIX versus CX, depicted in the bottom right panel, is indicative of the magnitude of the induced distortions in the VIX. After 13:00, there is a downward drift in the VIX strike range, followed by a period of more pronounced instability, resulting in a downward bias in VIX of up to 5% during the crash phase, succeeded by wild oscillatory upward swings, producing periodic 10 to 15% overvaluation thereafter. Thus, the VIX is plagued by large artificial shifts over this turbulent episode. Just as the market conditions deteriorated, the precision of the "fear gauge" evaporated.

In summary, the high-frequency VIX series was severely downward biased during the height of the flash crash due to a collapse of the liquidity in the option market. Equally remarkably, the VIX becomes strongly upward biased shortly thereafter, as the effective strike range expanded well beyond the levels observed prior to the crash. The latter likely reflects a partial restoration of confidence among market makers along with residual customer demand for downside protection during the still volatile aftermath of the crash. This induces the VIX index, erroneously, to reach a maximum for the day at around 14:30, about 45 minutes after the market trough.

We conclude by quantifying these effects through the correlations between the volatility indices

[29] The S&P 500 series stems from the CME Group. The E-mini futures contract is extraordinarily liquid and these futures prices are viewed as the most timely indicators of high-frequency developments in the S&P 500 index available.

Figure 11: **VIX and Corridor Implied Volatility on May 6, 2010** The top panels depict the evolution of the volatility indices (in units given on the left vertical scale) and the S&P 500 futures (in units given on the right vertical scale) across the trading day. The middle left (right) panel provides the corresponding evolution for the VIX and RX1 ($CX^*$) series. The bottom left panel displays the effective strike range used to compute RX1, while the bottom right panel conveys the relative discrepancy between the RX1 and $CX^*$ across the trading day.

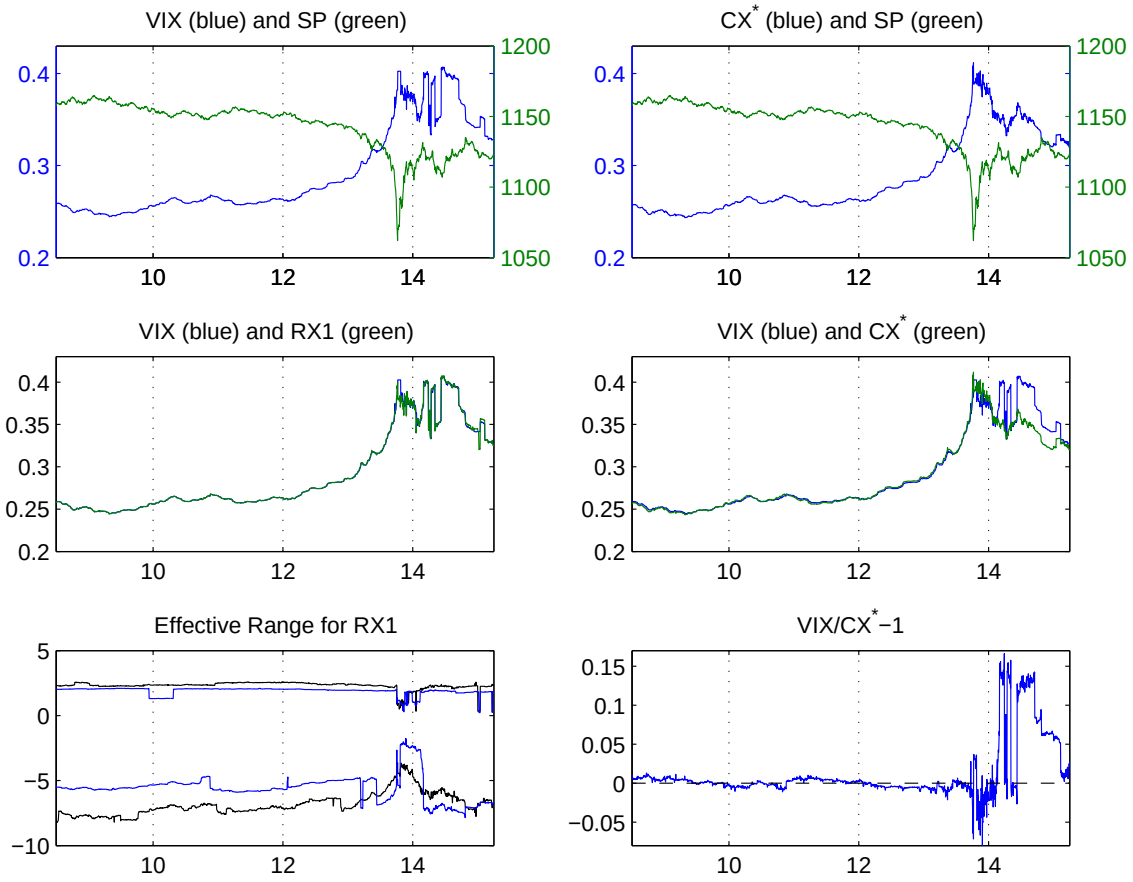


and the S&P 500 returns. Since the market disruptions also might have impaired the quality of the implied forward price computation on this day, we use the S&P 500 futures prices for computing the equity return.[30] Table 8 shows that all the volatility index returns were strongly negatively correlated with the equity-index returns before 13:30, as expected, even if the correlations were the smallest for VIX. These numbers are consistent with the correlations estimated for the full sample in Table 6 so, from this perspective, the market environment was not unusual up to this point. In contrast, after 13:30, except for CX, the correlations shrink dramatically. Hence, only CX retains a near invariant relation to the underlying S&P 500 index returns throughout the day. Moreover, the values for CX, based on observations covering only a few trading hours, continue to be in line with the ranges reported in the preceding section, corroborating the claim that we can produce sensible leverage coefficients from short intraday periods from high-frequency returns using our one-minute corridor volatility index. Furthermore, consistent with our prior evidence, the officially disseminated VIX series provides the worst coherence with the equity-index throughout the crisis period. The VIX index fails when it is

[30]The findings based on the implied forward prices were qualitatively similar, but the correlations were lower following the crash, consistent with the presence of relatively more significant distortions in the option market.

Table 8: **Volatility and S&P 500 Return Correlations, May 6, 2010.** The equity returns are based on the E-mini S&P 500 futures prices. The correlations are computed from one-minute log return observations.

| | $RX1$ | $RX2$ | $CX$ | $VIX^*$ |
|---|---|---|---|---|
| Before 13:30 | -0.70 | -0.72 | -0.77 | -0.55 |
| After 13:30 | -0.42 | -0.57 | -0.71 | -0.13 |
| All | -0.44 | -0.58 | -0.71 | -0.15 |

most needed – plagued by large idiosyncratic biases, arising endogenously as frictions grow more prominent and turbulence engulfs the markets.

# 7 Conclusions

We advocate the use of directly observable high-frequency implied volatility indices for exploring important features of the underlying return generating process. The approach is largely model-free and generally valid under standard no-arbitrage assumptions, as long as the risk-neutral volatility dynamics is governed by the current state of spot volatility. Although the volatility index differs from the spot volatility, it will display the identical instantaneous correlation with the underlying return process. Likewise, the volatility index shares important jump characteristics with the spot volatility process. They jump at the same time, and the sign and relative magnitude of volatility index jumps should generally be indicative of the corresponding features of spot volatility.

Upon analyzing the high-frequency behavior of the VIX index, we conclude that it is plagued by random idiosyncratic shifts in the strike range which render the measure incoherent. This induces artificial jumps in the measure and severely distorts its statistical properties. Consequently, it is not suited for applications that hinge on precise identification of high-frequency features. Instead, we develop the corridor implied volatility index, CX, which may be computed directly from real-time option quotes via a simple and transparent computational rule. The main change from the VIX is the determination of a different truncation point for when to exclude far OTM options from the calculation. By ensuring intertemporal consistency in this dimension, we obtain an index that covers economically equivalent parts of the strike range. Consequently, it is internally consistent, or coherent, and index values may be meaningfully compared across time, independently of the liquidity of the options market.

Using the CX index, we establish a number of new results regarding the nature of the high-frequency return dynamics. First of all, the volatility process displays fairly frequent jumps, but the individual jumps are much smaller than one may deduce from daily volatility measures, such as the daily VIX index. Moreover, the volatility jump distribution is largely symmetric and major volatility jumps are usually accompanied by price jumps of the opposite sign. We also document a significantly stronger negative spot return-volatility correlation than reported in the extant literature. In addition, the return-volatility asymmetry, or leverage effect, varies systematically with market conditions as it turns increasingly negative when uncertainty grows and return volatility is rising.

Our findings point towards a complex and dynamic interaction between the high-frequency return and volatility processes. An improved understanding of these features should assist in identifying the origin of the fat tails and asymmetries that are apparent in longer horizon return distributions. Moreover, it should help us construct models with a more realistic depiction of the return generating process and thus allow for pricing that incorporate risk compensation for the critical factors behind the high-frequency return dynamics and its propagation over longer horizons.

Finally, from a more policy-oriented perspective, our findings suggest it may be beneficial to secure

the continuous dissemination of a robust implied volatility measure to the marketplace. The weaknesses of the VIX stem from the cut-off rule applied for the strike range. Empirically, turbulent market conditions are tantamount to wide bid-ask spreads along with generally poor and rapidly shifting liquidity in the extreme part of the option strike range. A modification or supplement to the VIX index, exploiting the superior robustness and coherence properties of the corridor index construction, should help mitigate this problem. In particular, it would enhance the real-time flow of relevant pricing information during periods of market stress.

## References


Ait-Sahalia, Y., J. Fan & Y. Li. 2012. "The Leverage Effect Puzzle: Disentangling Sources of Bias at High Frequency." *Working Paper, Princeton University* .

Andersen, T.G. & O. Bondarenko. 2007. "Construction and Interpretation of Model-Free Implied Volatility." pp. 141–181. *Volatility as an Asset Class*, Edited by I. Nelken, London: Risk Books.

Andersen, T.G. & O. Bondarenko. 2010. "Dissecting the Pricing of Equity Index Volatility." *Working Paper, Northwestern University and University of Illinois at Chicago* .

Andersen, T.G. & T. Bollerslev. 1997. "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns." *Journal of Finance* 52:975–1005.

Andersen, T.G., T. Bollerslev & D. Dobrev. 2007. "No-Arbitrage Semi-martingale Restrictions for Continuous-Time Volatility Models Subject to Leverage Effects, Jumps and I.I.D. Noise: Theory and Testable Distributional Implications." *Journal of Econometrics* 138:125–80.

Andersen, T.G., T. Bollerslev & F.X. Diebold. 2010. "Parametric and Nonparametric Volatility Measurement.". L.P. Hansen and Y. Aït-Sahalia (eds.), *Handbook of Financial Econometrics*, North-Holland, 67-138.

Back, K. 1991. "Asset Pricing for General Processes." *Journal of Mathematical Economics* 20:371–395.

Bakshi, G., C. Cao & Z. Chen. 1997. "Empirical Performance of Alternative Option Pricing Models." *Journal of Finance* 52:527–566.

Bandi, F.M. & R. Reno. 2011. "Price and Volatility Co-Jumps." *Working Paper; Johns Hopkins University* .

Bandi, F.M. & R. Reno. 2012. "Time-Varying Leverage Effects." *Journal of Econometrics* 169:94–113.

Barndorff-Nielsen, O.E. & N. Shephard. 2001. "Non-Gaussian Ornstein-Uhlenbeck-based Models and some of their Uses in Financial Economics." *Journal of the Royal Statistical Society, Series B* 63, Part 2:167–241.

Barndorff-Nielsen, O.E. & N. Shephard. 2004. "Power and Bipower Variation with Stochastic Volatility and Jumps." *Journal of Financial Econometrics* 2:1–37.

Bates, D. 2000. "Post-'87 Crash Fears in S&P 500 Futures Options." *Journal of Econometrics* 94:181–238.

Bekaert, G. & G. Wu. 2000. "Asymmetric Volatility and Risk in Equity Markets." *The Review of Financial Studies* 13:1–42.

Black, F. 1976. "Asset Pricing for General Processes." *Proceedings of the 1976 Meetings of the American Statistical Association* pp. 171–181.

Bollerslev, T., J. Litvinova & G. Tauchen. 2006. "Leverage and Volatility Feedback Effects in High-Frequency Data." *Journal of Financial Econometrics* 4:353–384.

Bondarenko, O. 2000. "Recovering Risk-Neutral Densities: A New Nonparametric Approach." *Working Paper, University of Illinois, Chicago* .

Bondarenko, O. 2010. "Variance Trading and Market Price of Variance Risk." *Working Paper, University of Illinois, Chicago* .

Broadie, M., M. Chernov & M. Johannes. 2007. "Model Specification and Risk Premia: Evidence from Futures Options." *Journal of Finance* 62:1453–1490.

Campbell, J.Y. & L. Hentschel. 1992. "No News Is Good News: An Asymmetric Model of Changing Volatility in Stock Returns." *Journal of Financial Economics* 31:281–318.

Carr, P. & D.B. Madan. 1998. "Option Valuation Using the Fast Fourier Transform." *Journal of Computational Finance* 2:61–73.

Carr, P. & L. Wu. 2009. "Variance Risk Premiums." *The Review of Financial Studies* 22:1311–1341.

Christie, A.A. 1982. "The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects." *Journal of Financial Economics* 10:407–432.

Corsi, F. 2009. "A Simple Approximate Long-Memory Model of Realized Volatility." *The Journal of Financial Econometrics* 7:174–196.

Duffee, G.R. 1995. "Stock Returns and Volatility: A Firm-Level Analysis." *Journal of Financial Economics* 37:399–420.

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

Engle, R.F. & V. Ng. 1993. "Measuring and Testing the Impact of News in Volatility." *Journal of Finance* 48:1749–1778.

Eraker, B. 2004. "Do Equity Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices." *Journal of Finance* 59:1367–1403.

Eraker, B., M. Johannes & N. Polson. 2003. "The Impact of Jumps in Volatilityand Returns." *Journal of Finance* 58:1269–1300.

Fisher, R.A. 1921. "On the "Probable Error" of a Coefficient of Correlation Deduced from a Small Sample." *Metron* 1:332.

French, K.R., G.W. Schwert & R.F. Stambaugh. 1987. "Expected Stock Returns and Volatility." *Journal of Financial Economics* 19:3–29.

Glosten, L.R., R. Jagannathan & D.E. Runkle. 1993. "On the Relationship Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks." *Journal of Finance* 48:1779–1802.

Harvey, A.C. & N. Shephard. 1996. "The Estimation of an Asymmetric Stochastic Volatility Model for Asset Returns." *Journal of Business and Economic Statistics* 14:429–434.

Heston, S.L. 1993. "A Closed-Form Solution for Options with Stochastic Volatility and Applications to Bond and Currency Options." *The Review of Financial Studies* 6:327–347.

Huang, X. & G. Tauchen. 2005. "The Relative COntribution of Jumps to Total Price Variance." *Journal of Financial Econometrics* 3:456–499.

Ishida, I., M. McAleer & K. Oya. 2011. "Estimating the Leverage Effect of Continuous-Time Stochastic Volatility Models Using High-Frequency S&P 500 and VIX." *Working Paper; Osaka University* .

Jacod, J. & V. Todorov. 2010. "Do Price and Volatility Jump Together?" *The Annals of Applied Probability* 20:1425–1469.

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

Lee, S.S. & P.A. Mykland. 2007. "Jumps in Financial Markets: A new non-parametric test and Jump dynamics." *Review of Financial Studies* 20:2535–2563.

Mancini, Cecilia. 2009. "Non-parametric Threshold Estimation for Models with Stochastic Diffusion Coefficient and Jumps." *Scandinavian Journal of Statistics* 36:270–296.

Martin, I. 2012. "Simple Variance Swaps." *Working Paper; Stanford University* .

Nagel, S. 2012. "Evaporating Liquidity." *The Review of Financial Studies* 25:2005–2039.

Nelson, D.B. 1991. "Conditional Heteroskedasticity in Asset Returns: A New Approach." *Econometrica;* 59:347–370.

Todorov, V. 2010. "Variance Risk Premium Dynamics: The Role of Jumps." *The Review of Financial Studies* 23:345–383.

Todorov, V. & G. Tauchen. 2010. "Activity Signature Functions for High-Frequency Data Analysis." *Journal of Econometrics* 154:125–138.

Veraart, A.E.D. & L.A.M. Veraart. 2012. "Stochastic Volatility and Stochastic Leverage." *Annals of Finance* 8:205–233.

Wang, D.C. & P.A. Mykland. 2012. "The Estimation of Leverage Effect with High Frequency Data." *Working Paper; Department of Statistics, University of Chicago* .

# Appendix

## A Data and Data Filtering

### A.1 SPX Option Classes

We acquired the MDR (Market Data Retrieval) data for S&P 500 options from the CBOE subsidiary, Market Data Express (`http://www.marketdataexpress.com/`). The MDR data include tick-by-tick quotes and transactions throughout the trading day for all option classes issued by the CBOE on the S&P 500. Each option class is characterized by a letter code. We only consider options that the CBOE actually used in their computation of the VIX over our sample period, namely those in the SPB, SPQ, SPT, SPV, SPX, SPZ, SVP, SXB, SXM, SXY, SXZ, SYG, SYU, SYV and SZP categories. The latter are generally known as SPX equity options and they mature on the Saturday immediately following the third Friday of the expiration month (see `http://www.cboe.com/Products/EquityOptionSpecs.aspx`).

### A.2 Systemic Staleness in Option Quotes

By far, the most influential options for the computation of model-free implied volatility measures are the out-of-the-money (OTM) put and call options with strikes close to the at-the-money (ATM) strike which is given by the forward price. Hence, we actively monitor the liquidity of the ATM option as well as the set of twenty OTM put and twenty OTM call options closest to ATM for the two maturities exploited for the VIX computation. We label this group of options the "pivotal" ones.

Our systemic staleness filter flags episodes where there are no quote update among the pivotal options in the bid *or* the ask price at *either* of the two maturities for five minutes. Hence, this dummy variable equals one if all options within one of these four pivotal option groups have not been updated for five minutes and it is zero otherwise. When the flag is activated (equals unity), we classify the *entire* inactive period of five minutes or more as unreliable, and the volatility indices are not available (n.a.).

### A.3 Non-Convexity Filter

To preclude arbitrage opportunities, theoretical call and put prices must be monotonic and convex functions of the strike. In particular, the call prices must satisfy the following convexity restriction:

$$D_i \;=\; \frac{C(K_{i+1}) - C(K_i)}{K_{i+1} - K_i} - \frac{C(K_i) - C(K_{i-1})}{K_i - K_{i-1}} \geq 0\,,$$

and a similar restriction for the put prices. We obtain the option prices as the average of the bid and ask quotes and use the above restriction to identify "suspect" cross-section with apparent arbitrage violations, which could arise from recording errors, staleness, and other issues. Specifically, for each strike $K_i$, we compute the following measure of local non-convexity:

$$NC_i \;=\; -\min\{D_i, 0\}\,.$$

For low strikes ($K_i \leq F$), we compute $NC_i$ from OTM puts and, for high strikes ($K_i > F$), we use OTM calls. We then average $NC_i$ across all strikes to obtain the aggregate measure of non-convexity $NC$.

When $NC > 0.1$, we deem a cross-section unreliable and do not use it in our econometric analysis. For those cross-sections, the option prices indicate sizeable apparent arbitrage opportunities.

However, for some illustrations, we need to compute volatility indices even when the quality of data is poor ($NC > 0.1$). In those cases, prior to computing the volatility indices, we adjust option prices by running the so-called *Constrained Convex Regression* (CCR). This procedure has been implemented in Bondarenko (2000). Intuitively, CCR searches for the smallest (in the sense of least squares) perturbation of option prices that restores the no-arbitrage restrictions.

### A.4 Bounce-Back Filter

We detect situations where a volatility index experiences a large move, which is almost immediately offset by a jump of a similar magnitude, but in the opposite direction. Such behavior is usually indicative of a data error which impacts only a few recorded quotes. Such an error induces a jump in returns when it occurs, and then another one when the error is eliminated and the recorded quotes reverse to their appropriate level. Specifically, we label such an incident a " bounce-back" if the index jumps by at least 9 "sigmas" and either (i) more than 75% of this jump is reversed within one minute, or (ii) more than 80% of the jump is reversed within two minutes.

## B Robust Volatility Estimation

For a given volatility index, let the associated one-minute log return at time $t$ on trading day $d$ by $r_{dt}$. We assume this return may be represented as

$$r_{dt} \;=\; \sigma_d \, f_t \, z_{dt} \,,$$

where $\sigma_d$ is average volatility for trading day $d$, $f_t$ is a scaling factor which adjusts for the intraday volatility pattern, and $z_{dt}$ are i.i.d. random variables with zero mean and unit standard deviation, but not necessary normally distributed.

We estimate the daily volatilities $\{\sigma_d\}_1^{N_d}$ and intraday adjustment factors $\{f_t\}_1^{N_t}$ as follows. First, for fixed $t$, we compute the average squared return across all trading days,

$$\frac{1}{N_d} \sum_d^{N_d} r_{dt}^2 \;=\; f_t^2 \cdot \frac{1}{N_d} \sum_d^{N_d} \sigma_d^2 \, z_{dt}^2 \,.$$

Since $z_{dt}$ are i.i.d. with $E[z_{dt}] = 0$ and $\mathrm{Var}(z_{dt}) = 1$, we obtain,

$$E\left[ \frac{1}{N_d} \sum_d^{N_d} r_{dt}^2 \right] \;=\; f_t^2 \cdot \left( \frac{1}{N_d} \sum_d^{N_d} \sigma_d^2 \right) \;=\; V_{mean} \cdot f_t^2 \,,$$

where $V_{mean}$ is a constant independent of $t$. Therefore, the adjustment factor $f_t$ may be estimated as

$$f_t^2 \;=\; \frac{\frac{1}{N_d} \sum_d^{N_d} r_{dt}^2}{V_{mean}} \,,$$

where $V_{mean}$ now is determined from the condition that,

$$\sigma_d^2 \;=\; \frac{1}{N_t} E\left[ \sum_t^{N_t} r_{dt}^2 \right] \;=\; \frac{1}{N_t} \sum_t^{N_t} \sigma_d^2 \cdot f_t^2 \,,$$

or

$$\frac{1}{N_t} \sum_t^{N_t} f_t^2 \;=\; 1 \,. \tag{22}$$

The above approach provides an unbiased estimator of $f_t^2$. However, it might not be robust in the presence of extreme returns in the $r_{dt}$ series. Therefore, we implement a robustified version of the same approach, based on the sample median rather than the mean. Specifically,

$$E\left[ \mathrm{Median}_d \left\{ r_{dt}^2 \right\} \right] \;=\; f_t^2 \, \mathrm{Median}_d \left\{ \sigma_d^2 \right\} \;=\; V_{med} \cdot f_t^2 \,,$$

where $V_{med}$ again denotes a constant independent of $t$. The adjustment factor $f_t$ i then estimated as

$$f_t^2 \;=\; \frac{\mathrm{Median}_d \left\{ r_{dt}^2 \right\}}{V_{med}} \,,$$

subject to the constraint in (22). For additional robustness, the estimated intraday factors are smoothed by averaging them over 10-minute windows.[31] Armed with $\{f_t\}_1^{N_t}$, the daily volatility $\sigma_d$ can now be estimated as the volatility of the re-scaled returns, $u_{dt} = \frac{r_{dt}}{f_t}$. This could be done in many ways, but we focus on the robust estimator based on 5-to-95 percentile range. Specifically, we sort the re-scaled returns $u_{dt}$ for a given day $d$ and determine their 5- and 95-percentiles, $P(0.05)$ and $P(0.95)$. Then,

$$\sigma_d = \frac{P(0.95) - P(0.05)}{3.2898},$$

where the denominator equals the 5-to-95 percentile range of the standard normal random variable.[32]

Finally, when defining large index returns, we take into account both the daily measure of volatility $\sigma_d$ and the intraday adjustment factor $f_t$. That is, the move is deemed large if the absolute value of the ratio,

$$\left| \frac{r_{dt}}{\sigma_d f_t} \right| = \left| \frac{u_{dt}}{\sigma_d} \right|$$

exceeds a pre-specified threshold.

[31] For all volatility indices, the intraday factors range from about 0.6 to 2.0.

[32] We also experimented with other robust estimators of $\sigma_d$, and the results are quantitatively similar.