

Construction and Interpretation of Model-Free Implied Volatility*

Torben G. Andersen[†] and Oleg Bondarenko[‡]

August 2007

Abstract

The notion of model-free implied volatility (MFIV), constituting the basis for the highly publicized VIX volatility index, can be hard to measure with accuracy due to the lack of precise prices for options with strikes in the tails of the return distribution. This is reflected in practice as the VIX index is computed through a tail-truncation which renders it more compatible with the related concept of corridor implied volatility (CIV). We provide a comprehensive derivation of the CIV measure and relate it to MFIV under general assumptions. In addition, we price the various volatility contracts, and hence estimate the corresponding volatility measures, under the standard Black-Scholes model. Finally, we undertake the first empirical exploration of the CIV measures in the literature. Our results indicate that the measure can help us refine and systematize the information embedded in the derivatives markets. As such, the CIV measure may serve as a tool to facilitate empirical analysis of both volatility forecasting and volatility risk pricing across distinct future states of the world for diverse asset categories and time horizons.

JEL Classification: C53, G12, G13

Keywords: Model-Free Implied Volatility, Corridor Implied Volatility, Realized Volatility, VIX, Volatility Forecasting, Risk-Neutral Density

*The work of Andersen is supported by a grant from the NSF to NBER and by CREATES funded via the Danish National Research Foundation. The paper was initiated while Bondarenko was visiting Kellogg. Any errors remain our responsibility.

[†]Department of Finance, Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL; NBER and CREATES. Email: t-andersen@kellogg.northwestern.edu.

[‡]Department of Finance (MC 168), University of Illinois at Chicago, 601 S. Morgan St., Chicago, IL 60607. Phone: (312) 996-2362. Email: olegb@uic.edu

1 Introduction

Both market and academic interest in equity-index volatility measures has grown rapidly in recent years. The best known example is the publication of the so-called volatility index, or VIX, by the Chicago Board of Options Exchange (CBOE). Practitioners and academics alike have established that this index correlates significantly, not only with future equity market volatility, but also with global risk factors embedded in credit spreads or sovereign debt spreads. This explains the moniker of “global fear index” which has been attached to the VIX measure in the public press. Given the evident asset market implications, this is an area of growing interest for financial and macro economists. Moreover, public and over-the-counter markets have emerged to enable direct trading of volatility as measured by realized return volatility indicators over a prescribed future horizon for a host of different financial assets.

These developments have been facilitated by the development of a “model-free” implied volatility (MFIV) measure which, at least in principle, can be derived directly from a comprehensive cross-section of European put and call option prices with strikes spanning the full range of possible values for the underlying asset at option expiry. Recent research has confirmed that this pricing relationship is robust and remains approximately valid for a broad class of relevant return generating processes, including jump-diffusive semimartingales models. This contrasts sharply with the traditional Black-Scholes implied volatility (BSIV) measure which relies on a specific, and counterfactual, assumption on the return dynamics. The latter induces the well documented smile and smirk patterns in implied volatility across the strikes so the BSIV is a direct function of the particular option used for the computation. The VIX replaces the multitude of BSIV measures with a unique value obtained as a weighted average across all observed option prices with appropriate time to maturity and, as mentioned, it remains valid under general assumptions regarding the return dynamics.

Nonetheless, the requirements of theory for deriving the MFIV measure are not met by existing data so some approximations are inevitable. There are, in particular, practical limitations in terms of the existence of liquid options with strike prices covering the entire support of the return distribution. As such, robust computational procedures are crucial. On this dimension, the VIX has come in for criticism. For example, Jiang and Tian (2005b) find that the CBOE implementation introduces random noise as well as systematic errors into the index.

In order to assess the implications of this criticism and explore some conceptual issues concerning the construction of the volatility index it is useful to reflect on the notion it is intended to capture. The theoretical foundation for the MFIV, and thus the VIX, makes this transparent: the index aims to measure the expected integrated variance, or more generally return variation, over the coming month, evaluated under the so-called risk-neutral, or pricing (Q), measure. Since volatility is stochastic, the MFIV measure will typically differ from the expected return variation under the actual, or objective (P), measure. As such, the MFIV is not a pure volatility forecast for the underlying asset, but rather bundles this forecast with market pricing of the uncertainty surrounding the forecast. This implies that, in general, implied volatilities will include premiums compensating for the systematic risk associated with the exposure to equity-index volatility. Even so, all else equal, the volatility index will rise in response to a perceived increase in future (objective P -measure) volatility and vice versa. Consequently, the MFIV index should be strongly correlated with future realized volatility. In fact, since derivatives markets aggregate the views of many agents who make trading decisions based on current information stemming from diverse sources, including but vastly exceeding the information contained in historical returns, many scholars deem implied volatility forecasts

superior to other predictors generated by alternate methods, e.g., univariate time series models, on purely a priori grounds. Of course, this issue has been explored from different angles in the literature, but it has proven difficult to reach a firm conclusion due to the relatively short time span of reliable option prices, the large degree of noise present in standard measures of ex-post realized return volatility and, most importantly, the unobserved nature of the risk premiums embedded in implied volatilities.

It is impossible to decisively establish superiority of a given implied volatility index relative to others because any observed implied volatility level may be rationalized by appeal to an embedded unobserved premium. Nonetheless, there are a number of dimensions on which such indices may be assessed. First, consistent pricing of options of a given maturity across all possible strikes induces a risk-neutral density which must satisfy certain regularity conditions to exclude arbitrage opportunities. A primary requirement is that the risk-neutral density is strictly positive for all possible future values of the underlying asset. From this perspective, a striking feature of the computation of the VIX is that the CBOE truncates the tails of the return distribution at the point where no reliable option prices with corresponding strikes can be inferred. Moreover, the extent to which reliable option prices are available in these tail regions differs across trading days so the severity of the truncation varies stochastically over time. One sensible alternative is to apply a more theoretically coherent technique of extending the risk-neutral density into the tails, and we do so later in the chapter. However, any such procedure must rely on partially unverifiable assumptions, given the inherent data limitations. This will inevitably introduce a degree of random noise into the measure. Hence, it is intuitively appealing to focus on measures computed only over regions of the risk-neutral density where it may be inferred in a reliable fashion. Of course, such a truncated implied volatility measure should not be seen as representing the full MFIV but rather a deliberately down-scaled version of the latter. In fact, this interpretation can be formalized rigorously as the construction may be viewed as a variant of the so-called model-free “corridor implied volatility” (CIV) measure, briefly discussed by Carr and Madan (1998). Hence, it may be appropriate to view the VIX as an (imperfect) CIV index rather than a MFIV measure. Since the notion of CIV is not widely known and never, as far as we know, has been explored in practice, we provide a detailed exposition of the concept, linking it theoretically and empirically to the corresponding MFIV, VIX, and BSIV measures.

Secondly, we compare the forecast performance of alternative implied volatility measures. Although there is no requirement that a superior implied volatility measure is also a superior predictor of future volatility, inconsistently constructed or excessively noisy measures will tend to display poor coherence with the underlying market volatility movements and perhaps even contain predictable forecast errors which may be eliminated through alternative constructions of the measures. Hence, relative predictive ability can serve as an indirect indicator of the quality of the measures. Of course, it is also of independent interest to establish which implied volatility measures correlate most strongly with the underlying asset volatility and thus provide useful guidance for volatility prediction. We facilitate efficient inference regarding predictive performance by obtaining accurate measurements of the underlying ex-post realized volatility (RV). We rely on recently developed techniques for constructing realized volatility measures from high-frequency intraday return series on the underlying asset for this purpose.

Thirdly, we investigate the statistical properties of each candidate implied volatility series relative to relevant historical realized volatility measures. This sheds additional light on issues concerning the presence of risk premiums, systematic forecast biases and the presence of noise

in the series.

We study the above issues in an empirical setting where we can obtain relatively precise measures of the underlying return variation and we have access to high quality options data so that we may construct alternative implied volatility measures, both of the model-free and the Black-Scholes variety, with good accuracy. The market setting associated with the published VIX measure is not convenient in this regard due to the underlying being a cash equity index. The index has five hundred underlying stocks whose prices are never observed simultaneously so the resulting index is plagued by well-established lead-lag or non-synchronicity biases at high frequencies. This has implications for our ability to measure the underlying realized volatility with precision from the intraday index returns. Moreover, the corresponding options data not readily available. Such issues are alleviated greatly by instead using the S&P 500 futures market and the associated options which are traded on the Chicago Mercantile Exchange (CME). At the same time, the cash and futures return volatility are intimately linked so that a volatility measure for one market should serve as a good proxy for the volatility of the other.

In summary, this chapter argues that the VIX index is closely related to the concept of corridor implied volatility (CIV). We then relate different CIV measures, distinguished by corridor width, to the ideal MFIV, the BSIV and a couple of historical RV measures. In this first empirical study of CIV we establish that broad corridor CIV measures are near substitutes for sensible empirical measures of MFIV while some narrower corridor width CIV measures tend to mimic the BSIV measure. We systematically document the volatility forecast properties of these measures and find, in contrast to some existing evidence, that the narrow corridor or BSIV measures are more useful predictors of future volatility than the broad corridor, MFIV or VIX measures. The statistical properties of the various measures are consistent with the interpretation that the broad corridor implied volatility measures embed large and time-varying risk premiums which encroach on their usefulness as direct indicators of future volatility. Importantly, these findings should not be taken as a criticism of the MFIV concept. Instead, it points to practical empirical implications of the dual features of market based measures, namely as vehicles that simultaneously provide forecasts of future volatility and price the risk associated with this expected future (stochastic) return variation.

The paper unfolds as follows. Section 2 provides the basic theoretical exposition of the model-free, barrier and corridor implied volatility measures and their relationship to the Black-Scholes measure. In that context, we also provide as explicit comparison of the Black-Scholes prices for certain variance contracts with those prevailing in the market place. It serves to highlight the inadequacy of the Black-Scholes setting for understanding the market pricing of these newly developed variance products. In addition, we review the concept of realized return volatility which is used both as the ex-post measure of actual realized volatility and as ex-ante volatility forecast indicators. Section 3 describes the origin of our data and some details of the data cleaning and construction. Section 4 reports on the empirical results. We first provide some descriptive statistics to convey the basic behavior of the new corridor volatility measures relative to the more traditional volatility series. We then explore both the in-sample and out-of-sample performance of the various volatility measures as predictors of future volatility. Finally, Section 5 provides concluding remarks.

2 Theoretical Background

This section provides the theoretical foundation for the volatility concepts and measurements explored in our analysis. We begin with a formal introduction of the concept of barrier and corridor variance contracts. We next review the basic features of the so-called realized volatility measures which are used to obtain relatively accurate measures of the actual (ex-post) return variation of the underlying asset. Finally, we review the implementation procedures we adopt in order to convert the various alternative volatility concepts into practical measures amenable for empirical analysis.

2.1 Barrier Variance and Corridor Variance Contracts

Throughout this section we fix the current time at $t = 0$ and we consider only contracts which pay off at a future fixed date T . For $0 \leq t \leq T$, F_t denotes the time t value of the S&P 500 futures contract expiring at date T' where $T \leq T'$. Moreover, the prices of European put and call options with strike K and expiration date T are given by $P_t(K)$ and $C_t(K)$. To simplify the exposition, the risk-free rate is assumed to be zero.¹ In what follows, $k = K/F_t$ indicates the *strike-to-underlying ratio*, or *moneyness* of an options contract. Although moneyness, k , varies with the underlying price F_t , we suppress this time dependence for notational convenience. Thus, a put (call) is out-of-the-money (OTM) if $k < 1$ ($k > 1$), is at-the-money (ATM) if $k = 1$, and is in-the-money (ITM) if $k > 1$ ($k < 1$). We also use $\tau = T - t$ to denote time-to-maturity.

The option prices may be computed using the risk-neutral density (RND), denoted $h_t(F_T)$:

$$\begin{aligned} P_t(K) &= E_t^Q[(K - F_T)^+] = \int_0^\infty (K - F_T)^+ h_t(F_T) dF_T, \\ C_t(K) &= E_t^Q[(F_T - K)^+] = \int_0^\infty (F_T - K)^+ h_t(F_T) dF_T. \end{aligned}$$

The RND satisfies the relationship first exposted in Ross (1976), Breeden and Litzenberger (1978), and Banz and Miller (1978),

$$h_t(F_T) = \frac{\partial^2 P_t(K)}{\partial K^2} \Big|_{K=F_T} = \frac{\partial^2 C_t(K)}{\partial K^2} \Big|_{K=F_T} \quad (1)$$

Let $g(F_T)$ denote a general payoff at time T . The function $g(F_T)$ is assumed to have a finite second derivative which is continuous almost everywhere. Following Carr and Madan (1998) and Bakshi and Madan (2000), for any $x \geq 0$, $g(F_T)$ can be represented as:

$$g(F_T) = g(x) + g'(x)(F_T - x) + \int_0^x g''(K)(K - F_T)^+ dK + \int_x^\infty g''(K)(F_T - K)^+ dK. \quad (2)$$

By setting $x = F_0$ and taking expectations in equation (2), one obtains

$$E_0^Q [g(F_T)] = g(F_0) + \int_0^{F_0} g''(K)P_0(K)dK + \int_{F_0}^\infty g''(K)C_0(K)dK$$

¹In reality, of course, the risk-free rate is non-zero. However, in empirical tests below, we convert *spot* prices of options into *forward* prices (for delivery at time- T). To obtain forward prices, spot prices are multiplied by $e^{r_f(T-t)}$, where r_f is the risk-free rate over $[t, T]$. For example, the forward put price is $P_t(K) = e^{r_f(T-t)} P_t^s(K)$, where $P_t^s(K)$ is the spot put price. A similar approach has been used in, for example, Dumas, Fleming, and Whaley (1998).

$$= g(F_0) + \int_0^\infty g''(K)M_0(K)dK, \quad (3)$$

where $M_t(K)$ denotes the minimum of the put and call,

$$M_t(K) = \min(P_t(K), C_t(K)).$$

In other words, of the two plain vanilla options with strike K , $M_t(K)$ equals the price of the one that is currently out-of-the-money.

In the current setting the futures price process F_t is a martingale under the risk-neutral measure. Suppose it follows the general diffusion:

$$\frac{dF_t}{F_t} = \sigma_t dW_t, \quad (4)$$

where W_t is a standard Brownian motion and σ_t is a strictly positive, cadlag (stochastic) volatility process. Notice that we allow the volatility process to feature jump discontinuities. By Ito's Lemma,

$$g(F_T) = g(F_0) + \int_0^T g'(F_t)dF_t + \frac{1}{2} \int_0^T g''(F_t)F_t^2\sigma_t^2 dt, \quad (5)$$

which implies that

$$E_0^Q [g(F_T)] = g(F_0) + \frac{1}{2} E_0^Q \left[\int_0^T g''(F_t)F_t^2\sigma_t^2 dt \right]. \quad (6)$$

Combining equations (3) and (6), one finds that

$$E_0^Q \left[\int_0^T g''(F_t)F_t^2\sigma_t^2 dt \right] = 2 \int_0^\infty g''(K)M_0(K)dK, \quad (7)$$

It is convenient to define the down-barrier indicator function as follows,

$$I_t = I_t(B) = 1[F_t \leq B],$$

with B denoting the barrier. We now consider the contract with time T payoff equal to the (down-) *Barrier Integrated Variance*,

$$BIVAR_B(0, T) = \int_0^T \sigma_t^2 I_t(B) dt.$$

In other words, the contractual payment is given by the realized variance calculated only when the futures price lies below the barrier B . As B diverges to ∞ , the payoff approaches the standard integrated variance:

$$IVAR(0, T) = \int_0^T \sigma_t^2 dt.$$

Carr and Madan (1998) show how to synthesize the continuously-monitored barrier variance when the underlying process is continuous. The no-arbitrage value of the down-barrier variance contract can be derived from the relationship in (7). Suppose that the function $g(F_T)$ is chosen as

$$g(F_T) = g(F_T; B) = \left(-\ln \frac{F_T}{B} + \frac{F_T}{B} - 1 \right) I_T.$$

In the sequel we exploit the following properties of this $g(F_T)$ function,

- (a) it is equal to zero for all values of $F_T \geq B$,
- (b) its first derivative is continuous for all F_T ,

$$g'(F_T) = \left(-\frac{1}{F_T} + \frac{1}{B}\right) I_T,$$

- (c) its second derivative is continuous for all $F_T \neq B$,

$$g''(F_T) = \frac{1}{F_T^2} I_T.$$

The relationship in (7) then implies that the value of the barrier variance contract is

$$BVAR_0(B) = E_0^Q \left[\int_0^T \sigma_t^2 I_t dt \right] = 2 \int_0^B \frac{M_0(K)}{K^2} dK, \quad (8)$$

The square root of the above expression can be interpreted as the option-implied barrier volatility.

$$BIV_0(B) = \sqrt{2 \int_0^B \frac{M_0(K)}{K^2} dK}, \quad (9)$$

In the limiting case of $B = \infty$, the barrier implied volatility coincides with the so-called model-free implied volatility $MFIV_0$. The concept of the model-free implied volatility was developed in original work of Dupire (1993) and Neuberger (1994).² The concept is referred to as “model-free” because it does not rely on any particular parametric model, unlike the Black-Scholes implied volatility. CBOE uses this concept as the basis for its recently redesigned volatility index VIX.

The contract which pays the *corridor* variance can be constructed from two barrier variance contracts with different barriers. Let B_1 and B_2 denote the lower and the upper barriers and consider the contract with time T payoff,

$$CIVAR_{B_1, B_2}(0, T) = \int_0^T \sigma_t^2 I_t(B_1, B_2) dt,$$

where the indicator function $I_t(B_1, B_2)$ is defined as

$$I_t(B_1, B_2) = I_t = 1[B_1 \leq F_t \leq B_2].$$

In other words, this contract pays the *corridor variance*, or the variance calculated only when the futures price between the barriers B_1 and B_2 . The value of the corridor variance contract is

$$CVAR_0(B_1, B_2) = E_0^Q \left[\int_0^T \sigma_t^2 I_t dt \right] = 2 \int_{B_1}^{B_2} \frac{M_0(K)}{K^2} dK. \quad (10)$$

Carr and Madan (1998) also introduce the contract which pays future variance *along* a strike. This contract can be obtained as the limiting case of the corridor variance contract

²See also Carr and Madan (1998), Demeterfi, Derman, Kamal, and Zou (1999), and Britten-Jones and Neuberger (2000).

when the distance between the upper and lower barriers shrinks to zero and when the payoff is properly rescaled to have a non-negligible value.

$$SVAR_0(B) = \lim_{\Delta B \rightarrow 0} \frac{B}{\Delta B} CVAR_0(B, B + \Delta B) = 2 \frac{M_0(B)}{B}. \quad (11)$$

Note that we rescale the payoff by $\frac{B}{\Delta B}$ as opposed to just $\frac{1}{\Delta B}$ in Carr and Madan. This minor modification has the effect that, for Black-Scholes and most other canonical option models, the value of the along-strike variance contract depends only on the moneyness B/F_0 and not on the level of the underlying, F_0 , itself.

2.2 Barrier Variance Contracts under the Black-Scholes Model

In the Black-Scholes (1973) model, the instantaneous volatility is constant, $\sigma_t = \sigma$, and the value of the barrier variance can be computed in closed form as:

$$\begin{aligned} BVAR_0(B) &= 2 \left(N(y) (\sigma\sqrt{\tau}y - 1) + \frac{F_0}{B} N(y - \sigma\sqrt{\tau}) + \sigma\sqrt{\tau}n(y) - g(F_0) \right) \quad (12) \\ &= 2 \left(N(y) \left(-1 - \ln \frac{F_0}{B} + \frac{1}{2}\sigma^2\tau \right) + \frac{F_0}{B} N(y - \sigma\sqrt{\tau}) - \sigma\sqrt{\tau}n(y) + \left(1 - \frac{F_0}{B} + \ln \frac{F_0}{B} \right) I_0(B) \right), \end{aligned}$$

where

$$y = -\frac{\ln \frac{F_0}{B}}{\sigma\sqrt{\tau}} + \frac{1}{2}\sigma\sqrt{\tau},$$

and $n(\cdot)$ and $N(\cdot)$ denote the standard normal probability and cumulative density functions (pdf and cdf) respectively.

When $\sigma\sqrt{\tau}$ is small, the above expression can be approximately written as

$$BVAR_0(B) \approx \sigma^2\tau I_0(B).$$

Intuitively, if the current futures price is below the barrier ($F_0 < B$), then for small $\sigma\sqrt{\tau}$ it will remain there almost surely and the down-barrier variance is identical to the integrated variance. On the other hand, if the current futures price is above the barrier, the futures will remain above until maturity, almost surely, and the down-barrier variance is zero.

To provide some initial intuition, the top panel of Figure 1 plots the function

$$U_0(p) = \frac{BVAR_0(H_0^{-1}(p))}{BVAR_0(\infty)},$$

which is equal to the normalized barrier variance expressed in terms of the cumulative risk-neutral probability $p = H_0(B)$. By construction, the function $U_0(p)$ is monotonically increasing on $[0, 1]$ with $U_0(0) = 0$ and $U_0(1) = 1$. The figure assumes that $\tau = 21$ trading days and that σ is set to the average at-the-money implied volatility over the studied period.

For comparison, the bottom panel of Figure 1 plots the same function $U_0(p)$ computed from S&P 500 options, for a representative day in our dataset, 04/19/2000, when $\tau = 21$ trading days. The shape of the function $U_0(p)$ now differs dramatically from the Black-Scholes case, mainly reflecting the very fat left tail of the empirical RND. It is evident that the same features that induce systematic patterns in the Black-Scholes implied volatilities also will prevent the Black-Scholes model assumptions from delivering realistic market pricing of barrier variance contracts across the support of the RND.

2.3 Realized Volatility Measures

Given the futures price dynamics under the Q -measure specified in equation (4), the logarithmic futures price process under the actual P -measure will follow a semimartingale with the identical spot volatility process, σ_t . For short term (log-) price increments, the semimartingale property implies that the size of the innovation term is an order of magnitude larger than the size of the expected mean term. Hence, for high-frequency returns, the drift may be neglected. This line of reasoning is entirely general and provides a formal basis for the use of realized volatility as an ex-post measure of return variation. Assuming that we have $n + 1$ log-price observations available over the relevant measurement horizon, obtained at times $0 = t(0) < t(1) < \dots < t(n-1) < t(n) = T$, we may define realized volatility as the cumulative sum of squared returns,

$$RV_n(0, T) = \sum_{i=1}^n \left[\ln F_{t(i)} - \ln F_{t(i-1)} \right]^2.$$

Conditional on the observed price path over $[0, T]$, realized volatility provides an unbiased estimator of the underlying quadratic return variation, which is simply the integrated variance, $IVAR(0, T)$ in the current setting. Consequently, the conditional expectation at time $t = 0$ of the future quadratic return variation, denoted V_0 , will also equal the conditional expectation of future realized volatility, i.e.,

$$V_0 = E_0^P \left[\int_0^T \sigma_t^2 dt \right] = E_0^P [RV_n(0, T)]. \quad (13)$$

This relationship is critical as we cannot directly observe realizations of the integrated variance, while we can construct empirical measures of realized return volatility. Hence, the latter will serve as our empirical proxy for the former. Theory stipulates that we exploit as many returns in the computation as possible: the precision of the realized volatility measure improves as the sampling frequency increases and eventually, in the limit of continuous sampling, converges to the underlying integrated variance.³ In practice, the semimartingale property is violated at the highest sampling frequencies due to the presence of market microstructure noise in the recorded prices which may induce sizeable biases. Consequently, we follow the common procedure of aggregating five-minute squared futures returns over the course of the trading day to obtain a reasonably precise and unbiased empirical measure of the underlying return variation.⁴

We also need to measure the return variation during the overnight period when the futures market is closed. In accordance with the general properties for realized volatility, the measure will remain unbiased if we add the squared overnight return, obtained as the squared close-to-open logarithmic price change, to the measure obtained over the trading period. We denote (the square-root of) this RV measure, constructed from high-frequency data and the overnight return as indicated, RVH. Obviously, the absence of detailed information on the price evolution overnight has a detrimental impact on the overall precision of the measure but the effect is limited due to the comparatively low volatility associated with non-trading periods.

³This property is highlighted by, e.g., Andersen and Bollerslev (1998), Andersen, Bollerslev, Diebold and Labys (2001), Barndorff-Nielsen and Shephard (2001, 2002), and Meddahi (2002).

⁴Given the limited microstructure effects in the S&P 500 futures market, the analysis of Andersen, Bollerslev and Meddahi (2006) indicates that this should work well in practice. Alternatively, one may utilize even higher frequency intraday returns and exploit the robust procedures advocated recently by, e.g., Bandi and Russell (2005), Zhang, Mykland and Ait-Sahalia (2005), and Barndorff-Nielsen, Hansen, Lunde and Shephard (2006)

Even though the use of high-frequency return based realized volatility measures has become widespread in recent years, the reliance on daily squared returns for computation of realized monthly return volatility remains common. In order to facilitate comparison across studies and obtain a direct indication of the practical advantage of using the high-frequency based measure, we also include a realized volatility measure computed as (the square-root of) the cumulated daily squared returns, and denoted RVD, in the subsequent empirical analysis.

2.4 Construction of the Volatility Measures

In the empirical study, we compare and contrast properties of the interrelated volatility measures introduced above. These measures are constructed daily and assume a fixed horizon or time to maturity of $\tau = 1$ month (21 trading days).

We consider four corridor implied volatilities, denoted CIV1-CIV4, with barriers defined as fixed percentiles of the RND. Specifically, we compute the corridor implied volatility as

$$CIV_0(B_1, B_2) = \sqrt{2 \int_{B_1}^{B_2} \frac{M_0(K)}{K^2} dK}, \quad (14)$$

with the barriers chosen so that $B_1 = H_0^{-1}(p)$ and $B_2 = H_0^{-1}(1 - p)$ for $p = 0.25, 0.10, 0.05,$ and 0.025 for CIV1-CIV4, respectively. In other words, CIV1-CIV4 correspond to increasingly wider corridors, where the corridor for CIV1 covers the range from the 25th and 75th percentiles of the RND, the corridor for CIV2 covers the range from the 10th and 90th percentiles, and so on. The model-free implied volatility, MFIV, corresponds to the limiting case of $p = 0$. The two broad corridor CIV measures, CIV3 and CIV4 were chosen with an eye towards the largest width of the RND which may be estimated with precision for almost all trading days. They serve as potential proxies for the MFIV in that they capture the expected variation over very wide ranges of the RND but can be measured with better accuracy. The narrower corridor measures, CIV1 and CIV2, are included to highlight the different properties that arise from focusing on ranges of the RND which display a lower sensitivity to the variance risk premium.

To construct the corridor implied volatilities, we first estimate the RND via the *Positive Convolution Approximation* (PCA) method developed in Bondarenko (2003). The procedure exploits the relationship in equation (1) to infer the conditional RND $h_0(F_T)$ and directly addresses some important limitations of actual option data, namely that (a) options are only traded for a discrete set of strikes, as opposed to a continuum, (b) very low and very high strikes are usually unavailable, and (c) option prices contain substantial measurement errors stemming from nonsynchronous trading, price discreteness, and bid-ask bounce. The PCA method is fully nonparametric, guarantees arbitrage-free density estimates, controls against overfitting in a small sample setting, and has been shown to be accurate in simulations. In addition to the estimate for RND, the method also provides the input into computing the put pricing function $P_0(K)$, the risk-neutral cumulative density function $H_0(K)$, and the barrier variance function $BVAR_0(K)$ for arbitrary strikes. The latter allows us to compute the full set of corridor variance measures.

Other option-implied measures that we consider include the at-the-money Black-Scholes implied volatility (BSIV) and the CBOE's new volatility index (VIX). In theory, VIX should be very close to MFIV. However, as Jiang and Tian (2005b) point out the specific procedure adopted by the CBOE to compute the VIX index introduces several potential biases. They identify three types of approximation errors, namely (i) truncation errors due to the fact that

very low and very high strikes are not available in practice; (ii) discretization errors induced due to the numerical integration being implemented using a relatively coarse grid of available strikes; and (iii) errors arising from a Taylor series expansion approximation. In practice, the errors (ii)-(iii) are small and, in principle, can be rendered negligible through improved implementation. On the other hand, the errors stemming from (i) can be considerable and there is no simple solution.

Intuitively, the published VIX can be interpreted as a corridor implied volatility (CIV) measure with barriers set to the lowest and highest strikes that CBOE uses on a given day to compute the index. The fact that the barriers change stochastically from day to day, depending on the liquidity in certain segments of the options market, may induce systematic biases in the VIX.

Finally, besides option-implied measures, we also rely on realized volatilities, computed using either daily or high-frequency returns. The realized variance measure, RVD, is computed as the sum of 21 trading day close-to-close squared log-returns. The realized high-frequency data based variance measure, RVH, also covers 21 trading days, but now all the 5-minute intra-trading day and the overnight close-to-open squared log-returns within the month are cumulated.

3 Data

The full sample period is from January 1990 through December 2006. Our data stem from several sources. From the Chicago Mercantile Exchange (CME), we obtain daily prices of options on the S&P 500 futures and transactions data for the S&P 500 futures themselves. From the Chicago Board Options Exchange (CBOE), we obtain daily levels of the newly redesigned VIX index. Although the CBOE changed the methodology for calculating the VIX in September 2003, they have backdated the new index to 1990 using historical option prices. Finally, from the U.S. Federal Reserve, we obtain Treasury bill rates, which are used to proxy for the risk-free interest rate.

The S&P 500 futures have four different maturity months from the March quarterly cycle. The contract size is \$250 times S&P 500 futures price (before November 1997, the contract size was \$500 times S&P 500 futures price). On any trading day, the CME futures options are available for six maturity months: four months from the March quarterly cycle and two additional nearby months (“serial” options). The CME options expire on a third Friday of a contract month. However, the quarterly options expire at the market open, while the serial options expire at the market close. For the serial options, we measure time to maturity as the number of calendar days between the trade date and the expiration date. For the quarterly options, we use the number of calendar days remaining less one.

The option contract size is one S&P 500 futures. The minimum price movement is 0.05. The strikes are multiples of 5 for near-term months and multiples of 25 for far months. If at any time the S&P 500 futures contract trades through the highest or lowest strike available, additional strikes are usually introduced.

The CME options on the S&P 500 futures and options on the S&P 500 Index itself, traded on the CBOE, have been the focus of many empirical studies. For short maturities CME and CBOE option prices are virtually indistinguishable. Nevertheless, there are a number of practical advantages to using the CME options. First, as is well known, there is a 15-minute difference between the close of the CBOE markets and the NYSE, AMEX, and NASDAQ

markets, where the S&P 500 components are traded. This difference leads to non-synchronicity biases between the recorded closing prices of the options and the level of the Index. In contrast, the CME options and futures close at the same time (3:15 PM CST). Second, it is easier to hedge options using highly liquid futures as opposed to trading in the 500 individual stocks. On the CME, futures and futures options are traded in pits side by side. This arrangement facilitates hedging, arbitrage, and speculation. It also makes the market more efficient. In fact, even traders of the CBOE options usually hedge their positions with the CME futures. Third, an additional complication is that the S&P 500 index pays dividends. Because of this, to estimate the risk-neutral densities from CBOE options, one must make some assumptions about the index dividend stream. No such assumptions are needed for the CME futures options. A disadvantage of the CME options is their American-style feature. However, we conduct our empirical analysis in such a way that the effect of the early exercise feature is minimal.

For each trading day, we estimate the implied volatility measures CIV1-CIV4, MFIV and at-the-money BSIV. To obtain these values, we follow several steps, which are described in more detail in Appendix B. Briefly, the steps include 1) filtering out unreliable option data; 2) checking that the option prices satisfy the theoretical no-arbitrage restrictions; 3) inferring forward prices for European puts and calls; 4) estimating the RND for a continuum of strikes; and 5) estimating the implied volatilities. For illustration, Figure 2 depicts the BSIV, normalized option prices, and the risk-neutral pdf and cdf for a representative trading day in the sample.

The realized volatility measure RVD is computed from official daily closing prices on the S&P 500 futures, while the RVH measure is obtained from the last recorded transaction price within each five-minute interval over the trading period combined with the overnight change from the official closing price to the opening price the subsequent trading day.

4 Empirical Results

This section presents our empirical findings. We first review the basic statistical properties of the various volatility measures and then investigate their relative performance as predictors of the subsequent volatility of the underlying S&P 500 futures.

Related recent work with a focus on the properties of model-free implied volatility and its relation to general asset market dynamics include Andersen, Frederiksen and Staal (2006), Ang, Hodrick, Xing and Zhang (2006), Bakshi and Kapadia (2005), Bakshi and Madan (2006), Bliss and Panigirtzoglou (2004), Bollerslev, Gibson and Zhou (2007), Bollerslev and Zhou (2007), Bondarenko (2007), Carr and Wu (2004), Duan and Yeh (2007), Todorov (2007) and Wu (2004).

4.1 Basic Features of the Volatility Measures

The top two panels of Figure 3 depict the level and daily returns for the S&P 500 futures over our full sample period. The bottom panel plots the associated realized one-month return volatility series, RVH, along with the CBOE VIX index measure. As explained in the previous section, the VIX may be viewed as an indicator of future monthly volatility while RVH provides a measure of the actual realized volatility over that month. A couple of points are evident from the graph. First, there is good coherence between the VIX index and the ensuing market volatility. However, since RVH is recorded daily but represents monthly (future) volatility,

there is a great deal of induced serial correlation in this series.⁵ Hence, this feature must be interpreted with some care. Second, it is evident that the VIX series almost uniformly exceeds the subsequent realized volatility. This is consistent with earlier work establishing the presence of a substantial negative variance risk premium in the VIX index. In other words, investors are on average willing to pay a sizeable premium to acquire a positive exposure to future equity-index volatility. Of course, the CBOE VIX is computed on the basis of options written on the S&P 500 cash index while we compute the realized volatility from S&P 500 futures. This may involve a mismatch which could explain the large and persistent gap between the two volatility measures. In order to control for such effects we turn to an analysis of various model-free implied volatility measures computed directly from options on S&P 500 futures contracts which are compatible with the RVH series. We also include the VIX measure in the analysis to facilitate comparison with existing work.

Table 1 reports various summary statistics for nine volatility measures: RVD, RVH, VIX, BSIV, CIV1, CIV2, CIV3, CIV4, and MFIV over the full sample and two subsamples. Focusing on the top panel, we first note that the mean of the VIX is compatible with the level of the MFIV extracted from the CME futures options, and they both exceed the level of actual realized volatility of the underlying asset, whether measured by RVD or RVH, by a margin of more than 23 per cent (0.185 versus 0.150). Second, VIX is highly correlated with both the MFIV and the broadest corridor variance measure, CIV4, even if the VIX is slightly higher and a bit more persistent than the corresponding measures obtained from the CME option futures market. Since the VIX is constructed in a manner that, as argued earlier, effectively makes it a hybrid between a pure model-free implied volatility and a broad corridor implied volatility measure, it is reassuring that it does tend to mimic the behavior of these independently constructed series based on related derivatives data. This is further supported by the extremely high correlations between VIX, MFIV and CIV4 reported in Table 2. It suggests that we can study the qualitative links between the VIX and concepts like (regular) model-free implied volatility and corridor implied volatility through the corresponding features of the implied volatility measures obtained from the S&P 500 futures option market.

Moving to the empirical features of the corridor implied variance based measures, which have not been explored in the literature hitherto, we have, by construction, a monotonically increasing pattern in the mean level of corridor implied volatility as we progress from CIV1 to CIV4. In another manifestation of the large variance risk premium embedded in the options markets we observe that the mean of the CIV2 measure, covering only eighty per cent of the RND, also is significantly higher than the historical realized volatility of the underlying asset. Moreover, not surprisingly, the narrower corridor volatility measures are more stable than the corresponding broader measures in terms of higher serial correlation and lower sample standard deviation, skewness and kurtosis. Note also that the realized volatility measures, as expected, are more volatile than the implied measures. They have the highest standard deviation, skewness and kurtosis statistics of all the series and they have the lowest serial correlation at monthly and lower frequencies where the measurement overlap ceases to have an effect. Such discrepancies are, of course, typical when comparing series that represent expectations of future realizations versus the actual ex-post realizations. The latter embody both an expected component, highly correlated with the implied volatility measures, and an

⁵Consecutive daily observations on the future monthly realized volatility share twenty of the twenty-one trading day realized return variation measures that are cumulated to provide the monthly realized volatility measure. Thus, only at the twenty-one trading day (monthly) frequency are the series not mechanically correlated.

unpredictable innovation term. The presence of the second component will naturally render the realized volatility series comparatively erratic in nature.

The final volatility measure included in the analysis is the traditional ATM Black-Scholes measure, BSIV. Although this measure also incorporates a significant variance risk premium, it is an order of magnitude smaller than for the MFIV. In addition, it is noteworthy that BSIV is extremely highly correlated with the intermediate corridor implied volatilities, and especially CIV2. Nonetheless, it is more persistent than these mid-range corridor volatility measures and sports lower sample skewness and kurtosis. In fact, at monthly and lower frequencies only the CIV1 measure is more persistent than BSIV.

A last observation concerns the very strong degree of persistence in the volatility measures. The autocorrelation patterns decay extremely slowly and are well approximated by a hyperbolic shape in all instances, thus lending support to the hypothesis that the volatility process has long memory components. Since this issue is not of direct relevance for our current study we abstain from additional analysis of these features.

The summary statistics for the subsamples are in line with those discussed above. As is also evident from the bottom panel of Figure 3, the average volatility is higher in the second subsample than the first, thus rendering the (right-skewed) volatility outliers less influential in the computation of the skewness and kurtosis for the second subsample. These separate subsamples play a pivotal role in the forecast analysis below.⁶

In summary, the various implied volatility measures are all highly correlated although clearly not identical. Figure 4 displays the time-variation of the various corridor implied volatility measures relative to MFIV and BSIV. The comparatively stronger coherence among the MFIV and the broader measures CIV3 and CIV4 as well as the very high correlation between BSIV and CIV2 is clearly visible in the top and bottom panels respectively. Another common feature across the implied volatility measures is that they all embed a sizeable variance risk premium. Concurrent research, e.g., Todorov (2007), concludes that the variance risk premium in the VIX tends to grow more than proportionally with the level of underlying volatility. The stationary and mean-reverting nature of realized volatility and the pronounced mean reversion evident from the plots in Figure 4 then suggest that the variance risk premium associated with the tails of the RND are particularly volatile. If this line of reasoning has a degree of validity, MFIV may be a poorer predictor of future underlying volatility than alternative, less premium sensitive, implied volatility measures such as BSIV and CIV1. On the other hand, recent work by Jiang and Tian (2005a) reach the exact opposite conclusion as they find their MFIV measure to dominate all other volatility forecasts, including the BSIV. Since our understanding of the variance risk premium dynamics and its manifestation across the different implied volatility measures is related to this issue, we now turn to a direct investigation of the predictive ability of the various volatility indicators. Of course, the findings are also highly relevant for the more general volatility forecasting literature.

4.2 The Relative Forecast Performance of Implied Volatility Measures

There is no simple way to rank alternative volatility forecasts as the relative performance generally will depend on the intended use of the predictions. In other words, different loss functions applied to a given forecast error distribution will typically not provide the identical

⁶The correlations among the volatility measures across the two subsamples were nearly identical to the corresponding statistics for the full sample reported in Table 2. Hence, for brevity, we do not report these correlation statistics for the shorter sample periods.

ranking. Given this fundamental problem, we resort to the well-known root-mean squared forecast error (RMSE) criterion which is widely used, simple to implement and equipped with convenient statistical properties.

An intuitively appealing starting point from which to explore the predictive ability of the different candidate volatility measures is to include each separately within an in-sample forecast regression, also known as a Mincer-Zarnowitz regression. Letting the ex-post realized volatility measure for month $t + 1$ be given by y_{t+1} and the volatility predictor j , among a set of J candidate predictors, be denoted $x_{j,t}$, $j = 1, \dots, J$, these univariate regressions take the form,

$$y_{t+1} = \alpha_j + \beta_j x_{j,t} + u_{j,t+1}, \quad (15)$$

where unbiased forecasts are subject to the constraint on the regression coefficients that $\alpha_j = 0$ and $\beta_j = 1$. Moreover, the regression R^2 captures the degree of variation in the ex-post realized volatility explained by the forecast. Likewise, one may explore whether one predictor candidate, x_j subsumes another, x_k , by including both in an encompassing regression of the form,

$$y_{t+1} = \alpha_{jk} + \beta_j x_{j,t} + \beta_k x_{k,t} + u_{j,t+1}, \quad (16)$$

where there is support for the hypothesis that x_j subsumes the information content in x_k if $\beta_j > 0$ and $\beta_k = 0$.

It is clear from our prior findings that the implied volatility measures generally will not be unbiased as they embed a sizeable premium related to equity market volatility risk. Nonetheless, they may well correlate strongly with future volatility and thus serve as useful indicators for prediction. Instead of imposing a priori restrictions on the character of the risk premiums, and thus the ensuing forecast biases, we allow the regressions to provide optimal (in-sample) regression coefficients for (linearly) transforming the specific volatility measures into forecasts of future realized volatility for the underlying asset. One immediate concern is that this may induce a small-sample bias, or tendency for over-fitting, given the relatively short history of monthly volatility forecasts, the strong persistence in the realized and implied volatility measures and even the possibility of regime shifts. To assess robustness against such concerns, we split the full sample into an in-sample period of estimation, where the OLS regression coefficients are obtained, and an out-of-sample period in which we keep the regression coefficients fixed and use them to construct monthly volatility forecasts on the basis of the subsequent implied volatility measures constructed from available options data. In other words, no estimation or recalibration is performed in the out-of-sample period. If the in-sample results are reliable, there should be good coherence in the ranking of performance across the two distinct sample periods. We select the initial ten years, covering 1990-1999, as the in-sample period and then explore robustness of the results in the seven year out-of-sample period comprising the years 2000-2006. The separation into the estimation and out-of-sample forecast period is indicated by the vertical lines in Figure 3.

An additional complication is that the return variation concept often differs across studies depending on the intended application or research question. Most commonly, it is reported on an annual basis in units of standard deviation or volatility, as is standard for Black-Scholes implied volatility. This is also the convention adopted by the CBOE when publicizing the VIX. Finally, it has a rationale within the stochastic volatility literature, as the ATM BSIV is approximately linear in the expected integrated volatility of the underlying asset up to expiry. Given the focus on return volatility in the literature, the majority of our exposition concentrates on forecast results for this concept of future monthly realized volatility. However, the theory

for model-free implied volatilities is developed for the return variance so a number of relations involving MFIV are more naturally couched in terms of variances. Finally, the small sample properties of the predictive regressions are decidedly better when return variation is measured in log volatilities as this eliminates the main positive outliers and renders the various series close to Gaussian. As a consequence, many prior studies focus on this metric as well. For robustness and compatibility with earlier work, we therefore provide supplementary results for predictive regressions targeting the future monthly log return volatility as well as the future monthly return variance. Since our empirical results turn out to be fully consistent across these settings, we only provide a brief summary of these additional findings.

Results for the full set of predictive regressions along with selected encompassing regressions for future return volatility are given in Table 3, while the supplementary results for log return volatility and return variance are reported in Tables 4-5. Focusing on the columns on the left side of Table 3, we first note that the VIX provides the worst in-sample fit among all candidate implied volatility measures although it comes close to the theoretically related model-free and broad corridor implied measures, MFIV and CIV4. The slightly narrower corridor measure CIV3 seems to perform a bit better while the most narrow corridor measures CIV1 and CIV2 along with BSIV perform the best. However, overall the in-sample fit does not differ dramatically across the implied volatility measures so a more definitive conclusion must await the findings from the, in practice, more challenging out-of-sample evidence. As expected all the implied volatility forecasts vastly outperform the benchmark consisting of the lagged realized (historical) volatility measures, RVD and RVH, although the latter explain a fairly impressive 54% of the variation in subsequent monthly volatility. The fact that both the historical volatility indicators have a slope coefficient below unity and a significant positive intercept reflects the mean-reverting character of realized volatility. More elaborate autoregressive time series models, estimated from a long history of daily realized volatility measures, tend to perform well, see, e.g., Andersen, Bollerslev, Diebold and Labys (2003), but it is beyond the scope of the present chapter to pursue alternative time series volatility forecast procedures. We simply note that timely high-frequency conditioning information is lost when aggregating the daily realized volatility series into monthly measures. As such, the historical volatility series, RVD and RVH, are mainly included to provide a simple and intuitive lower bound on the degree of volatility predictability over the sample.

The encompassing regressions provide additional insights. First, to the extent that lagged realized volatility conveys relevant information it is captured entirely by the high-frequency based measure RVH rather than the daily return based measure, RVD. The former has the superior in-sample fit, the encompassing regression including both measures adds no significant explanatory power to what is provided by RVH alone and the regression coefficient associated with RVD is now negative and insignificant. This is consistent with the findings in the literature that realized return variability is measured much more accurately via high-frequency return observations. This also supports our use of RVH as the ex-post measure of monthly realized return variation for the dependent variable on the left hand side of the predictive regressions. Second, note that the extremely strong correlation among the implied volatility measures suggests that little can be gained by exploiting two of these simultaneously in the forecast regression setting. This turns out to be true. Results for a representative scenario using both CIV1 and VIX are reported in the table. This pair constitutes a serious candidate for constructing a combined measure as the two displays the lowest correlation among our implied volatility indicators. Nonetheless, the improvement of the in-sample forecast performance is

slight and the coefficient on VIX is insignificant. In contrast, if the RVH measure is combined with CIV1 we obtain a larger increase in R^2 and both coefficients are significant. Hence, even if RVH in isolation represents a poorer volatility forecast than VIX, it adds more useful forecast information to the corridor variance measure than the VIX. Hence, the in-sample evidence suggests that, in terms of predictive relevance, the information embedded in the VIX is subsumed by our corridor variance measures while the historical volatility contains useful independent information.

We now turn to the pivotal out-of-sample evidence reported in the central and right columns of Table 3. The overall measure of forecast performance is the percentage (normalized) RMSE. If we let \hat{y}_t denote a forecast for y_t , it is formally defined as,

$$\text{RMSE} = \frac{\sqrt{E[(\hat{y} - y)^2]}}{\sqrt{E[y^2]}} \cdot 100.$$

The middle column labeled “All days” covers the full out-of-sample period. The findings are consistent with the in-sample results but the relative ranking is even more evident. The narrow corridor implied volatility measures and the BSIV continue to provide superior forecasts, with CIV1 now sporting the best overall performance. Forecast precision deteriorates monotonically as we move to CIV3, CIV4, MFIV and VIX. As before, the lagged monthly realized volatilities perform significantly worse than all implied volatility measures.

The last three columns of the table document performance over subsamples obtained by sorting the monthly forecasts in ascending order of at-the-money Black-Scholes implied volatility (BSIV). Hence, results for a third of the monthly forecasts, corresponding to the lowest BSIV measures, are provided in the “Low volatility” column, results for the next third are in the “Medium volatility” column, and for the last third, associated with high BSIV measures, in the last column. First, we note that the historical volatility series perform particularly poorly in the extreme segments of the volatility distribution. It is apparent that the long backward-looking nature of the historical measures constitutes a major disadvantage in terms of providing timely signals concerning the current (and likely future) level of return volatility which is most problematic whenever the current volatility level is unusually high or low. Second, all measures perform comparably in the low volatility regime in terms of normalized RMSE. The slightly better performance of VIX than the other implied volatilities is likely due to idiosyncratic sampling variation. In any case, the VIX and MFIV are the clear losers among the implied volatility measures for the higher volatility scenarios. In fact, the ranking across these two regimes is consistent with the overall findings as top performers listed in order are CIV1, BSIV, CIV2, CIV3, CIV4 and finally the two full-fledged model-free implied measures, MFIV and VIX.

The encompassing regressions largely confirm the observations drawn from the univariate out-of-sample predictive regressions. Since the coefficients are fixed at values obtained from the estimation sample, it is now feasible, and indeed common, for the combined forecasts to underperform the univariate predictors. In particular, the VIX again adds no value beyond what is captured by CIV1 as the combined forecasts fare worse than the CIV1 forecasts. Moreover, the RVH continues to supplement CIV1 better than the VIX although the indicated improvement, relative to forecasts generated by CIV1 alone, is slight.

The two supplementary set of results, based on the log volatility and variance measures of the realized return variation and the corresponding measures for the predictor variables, serve to underscore the robustness of the main qualitative findings. This is especially evident

in Table 4 where, remarkably, the ranking of the top five log-volatility forecasts is identical across all three volatility sorted subsamples and also consistent with the ranking obtained for the volatility measures above, while CIV4, MFIV and VIX always perform worse and the historical realized log-volatility measures remain at the bottom. Moreover, the combination of CIV1 and RVH within one encompassing regression now appears even more successful in adding explanatory power beyond what is captured by CIV1 alone, while VIX remains subsumed by CIV1 in this setting. In contrast, the evidence is slightly less clear cut in the more noise laden regression environment associated with Table 5. Even so, the out-of-sample findings still produce the same ranking among the implied forecasts as before.

One obvious concern with this comparative predictive analysis is that the conclusions may be driven by idiosyncratic features in the in-sample or out-of-sample periods. In particular, the forecast precision of some implied measures such as the VIX may be more sensitive to the discrepancy in average volatility across the subsamples than others. As a final robustness check, we reverse the roles of the estimation and prediction sample. Hence, we run the in-sample regressions for the seven year period 2000-2006, fix the regression coefficients at the point estimates obtained for this period, and then use historical realized volatilities and observed implied volatilities over 1990-1999 to forecast realized volatility in this “out-of-sample” period. The results provide strong confirmatory evidence. The ranking of forecast performance is identical to the one obtained earlier. Furthermore, this ranking is uniform across the volatility sorted subsamples. Moreover, in all instances, the combined forecast exploiting both CIV1 and RVH is superior to using the best single implied volatility predictor, CIV1, alone. Finally, the forecast information provided by the VIX measure continues to be subsumed by the information in CIV1.

5 Conclusion

This paper provides the first empirical study of corridor implied volatility, or CIV, measure in the literature. We find that broad corridor CIV measures serve as good substitutes for model-free implied volatility, or MFIV, with the advantage that CIV can be measured with better precision due to the lack of liquid options quotes in the tails of the risk-neutral distribution (RND). On the other hand, narrow corridor CIV measures are more closely related to the concept of (at-the-money) Black-Scholes implied volatility (BSIV). As such, they seem less sensitive to time variation in the market volatility risk premium which renders their time series behavior relatively more stable and allows them to serve as a superior gauge for the future volatility of the underlying asset returns, not only relative to MFIV but also BSIV. Hence, our findings suggest that the best possible market-based implied volatility measure for volatility prediction may take the form of a CIV measure. Even so, there are indications that historical realized volatility contains additional information for future volatility. It is an intriguing research question to determine how best to combine the implied and historically observed volatility measures for forecast purposes.

One general implication is that the MFIV measure should be interpreted strictly for what it seeks to represent, namely the market price of volatility exposure consistent with observed option prices. As such, it is a theoretically superior construction to the BSIV measure and it serves as a natural gauge for the pricing of broad asset categories with a strong exposure to general market risk. In contrast, in terms of direct indicators for future volatility, it must be recognized that MFIV combines volatility forecasting with pricing of the risk associated with

volatility. Consequently, even if MFIV provides a pure market based measure of the future return variation, the strong variation in the market pricing of volatility risk renders the linkage to the volatility process of the underlying asset tenuous. Nonetheless, our finding that the predictive content of the MFIV, and the VIX, is fully subsumed by the information conveyed by a narrow corridor CIV measure, or the BSIV, is new to the literature and contrary to some previous findings. However, our conclusion is based on a much longer time series and more carefully constructed volatility measures than existing studies of the relative predictive ability of VIX, or MFIV, and BSIV. The robustness of our finding is collaborated by the striking monotone improvement in forecast performance of the CIV measures as the corridor width is narrowed as well as the consistent results obtained across differently sorted subsamples.

In addition, we argue that the lack of liquid options in the tails of the RND induces inevitable measurement errors in MFIV which may be mitigated by restricting attention to a related CIV measure. Given how the VIX is actually constructed, this is in fact a fairly accurate descriptive of current practice. As such, explicit acknowledgment of the need and desirability of using a corridor measure in lieu of the full-fledged MFIV may motivate a modified implementation strategy which computes the measure on the basis of a predefined range of the RND rather than via a more error prone random truncation procedure.

A constructive message of our study is that judiciously selected CIV measures can be exploited in a theoretically coherent and empirically tractable manner to further refine the information embedded in the derivatives markets. For example, combining complementary CIV measures should enable us to detect variation in the pricing of equity volatility risk across distinct future states of the world in a timely fashion. This will facilitate more detailed studies of the interaction between the pricing of equity volatility and the conditions of related financial markets. In particular, this may shed some new light on the well documented, but poorly understood, linkages between the VIX index, the overall functioning and liquidity of the financial system and the pricing in global equity, credit and debt markets.

Finally, as long as the requisite option markets are sufficiently active, these tools are applicable across any asset class and horizon, irrespective of whether an official volatility index is being compiled and released or not. Hence, the notions of model-free and corridor implied volatility are not tied to equity-index volatility pricing and forecasting over a monthly horizon, but rather provide useful tools for quantifying and interpreting corresponding dynamic market features across diverse asset categories and maturities.

Appendix

A Derivation of Value of the Barrier Variance under Black-Scholes

One way to derive (12) is to write

$$\begin{aligned} E_0^Q [g(F_T)] &= \int_0^\infty g(K)h_0(K)dK = - \int_0^\infty g'(K)H_0(K)dK \\ &= \int_0^\infty \left(\frac{1}{K} - \frac{1}{B} \right) I_0 H_0(K)dK = \int_0^B \frac{H_0(K)}{K} dK - \frac{P_0(B)}{B}, \end{aligned}$$

where $H_t(K)$ is the risk-neutral cumulative density function. For the Black-Scholes model,

$$H_0(K) = N(z), \quad z = z(K) := -\frac{\ln \frac{F_0}{K}}{\sigma\sqrt{\tau}} + \frac{1}{2}\sigma\sqrt{\tau}.$$

The formula in (12) now obtains by noting that

$$\int_0^B \frac{H_0(K)}{K} dK = \int_0^y N(z)\sigma\sqrt{\tau}dz = \sigma\sqrt{\tau}(yN(y) - n(y)),$$

and

$$\frac{P_0(K)}{B} = N(y) - \frac{F_0}{B}N(y - \sigma\sqrt{\tau}).$$

B Construction of Dataset

To construct our dataset, we follow several steps:

1. For both options and futures we use settlement prices. Settlement prices (as opposed to closing prices) do not suffer from nonsynchronous/stale trading of options and the bid-ask spreads. CME calculates settlement prices simultaneously for all options, based on their last bid and ask prices. Since these prices are used to determine daily margin requirements, they are carefully scrutinized by the exchange and closely watched by traders. As a result, settlement prices are less likely to suffer from recording errors and they rarely violate basic no-arbitrage restrictions. In contrast, closing prices are generally less reliable and less complete.

2. In the dataset, we match all puts and calls by trading date t , maturity T , and strike. For each pair (t, T) , we drop very low (high) strikes for which put (call) price is less than 0.1. To convert spot prices to forward prices, we approximate the risk-free rate r_f over $[t, T]$ by the rate of Tbills.

3. Because the CME options are American type, their prices $P_t^A(K)$ and $C_t^A(K)$ could be slightly higher than prices of the corresponding European options $P_t(K)$ and $C_t(K)$. The difference, however, is very small for short maturities that we focus on. This is particularly true for OTM an ATM options.⁷

⁷As shown in Whaley (1986), the early exercise premium increases with the level of the risk-free rate, volatility, time to maturity, and degree to which an option is in-the-money.

To infer prices of European options $P_t(K)$ and $C_t(K)$, we proceed as follows. First, we discard all ITM options. That is, we use put prices for $K/F_t \leq 1.00$ and call prices for $K/F_t \geq 1.00$. Prices of OTM and ATM options are both more reliable and less affected by the early exercise feature. Second, we correct American option prices $P_t^A(K)$ and $C_t^A(K)$ for the value of the early exercise feature by using Barone-Adesi and Whaley (1987) approximation.⁸ Third, we compute prices of ITM options through the put-call parity relationship

$$P_t(K) + F_t = C_t(K) + K.$$

4. We check option prices for violations of the no-arbitrage restrictions. To preclude arbitrage opportunities, call and put prices must be monotonic and convex functions of the strike. In particular, the call pricing function $C_t(K)$ must satisfy

$$(a) \quad C_t(K) \geq (F_t - K)^+, \quad (b) \quad -1 \leq C_t'(K) \leq 0, \quad (c) \quad C_t''(K) \geq 0.$$

The corresponding conditions for the put pricing function $P_t(K)$ follow from put-call parity. When restrictions (a)-(c) are violated, we enforce them by running the so-called *Constrained Convex Regression* (CCR), see Bondarenko (2000). Intuitively, CCR searches for the smallest (in the sense of least squares) perturbation of option prices that restores the no-arbitrage restrictions. For most trading days, option settlement prices already satisfy the restrictions (a)-(c). Still, CCR is a useful procedure because it allows one to identify possible recording errors or typos.

5. For each pair (t, T) , we estimate RND using the *Positive Convolution Approximation* (PCA) procedure of Bondarenko (2000, 2003). Armed with RND, we may obtain the put and call pricing functions as well as the other fundamental objects used in our analysis.

⁸It is important to point out that this correction is always substantially smaller than typical bid-ask spreads. In particular, the correction generally does not exceed 0.2% of an option price.

References

- Andersen, T.G. and T. Bollerslev, 1998, "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review* 39, 885-905.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys, 2001, "The Distribution of Exchange Rate Volatility," *Journal of the American Statistical Association* 96, 42-55.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys, 2003, "Modeling and Forecasting Realized Volatility," *Econometrica* 71, 579-625.
- Andersen, T.G., T. Bollerslev and N. Meddahi, 2006, "Realized Volatility Forecasting and Market Microstructure Noise," Working Paper; Kellogg School, Northwestern University, Duke University and Imperial College.
- Andersen, T.G., P.H. Frederiksen and A. Staal, 2006, "The Information Content of Realized Volatility Forecast," Mimeo; Kellogg School of Management, Northwestern University.
- Ang, A., R.J. Hodrick, Y. Xing and X. Zhang (2006), "The Cross-Section of Volatility and Expected Returns," *Journal of Finance* 61, 259-299.
- Bakshi G.S. and N. Kapadia, 2003, "Delta-Hedged Gains and the Negative Market Volatility Risk Premium," *Review of Financial Studies* 16, 527-566.
- Bakshi, G.S., and D. Madan, 2000, "Spanning and Derivative-Security Valuation," *Journal of Financial Economics* 55, 205-238.
- Bakshi G.S. and D. Madan, 2006, "A Theory of Volatility Spreads," *Management Science* 52, 1945-1956.
- Bandi, F. and J. Russell, 2005, "Microstructure Noise, Realized Volatility, and Optimal Sampling," Working Paper; University of Chicago.
- Banz, R., and M. Miller, 1978, "Prices for State-Contingent Claims: Some Estimates and Applications," *Journal of Business* 51, 653-672.
- Barndorff-Nielsen, O.E., P.R. Hansen, A. Lunde and N. Shephard, 2006, "Designing Realised Kernels to Measure the Ex-Post Variation of Equity Prices in the Presence of Noise," Working Paper, Oxford University.
- Barndorff-Nielsen, O.E. and N. Shephard, 2001, "Non-Gaussian OU based Models and some of their uses in Financial Economics," *Journal of the Royal Statistical Society B* 63, 167-241.
- Barndorff-Nielsen, O.E. and N. Shephard, 2002, "Econometric Analysis of Realised Volatility and its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society B* 64, 253-280.
- Barone-Adesi, G., and R. Whaley, 1987, "Efficient Analytical Approximation of American Option Values," *Journal of Finance* 42, 301-320.
- Black, F., and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81, 637-654.

- Bliss, R.R. and N. Panigirtzoglou, 2004, "Option-Implied Risk Aversion Estimates," *Journal of Finance* 59, 407-446.
- Bollerslev, T., M. Gibson and H. Zhou (2007), "Dynamic Estimation of Volatility Risk Premia from Option-Implied and Realized Volatilities," Working Paper; Duke University and Federal Reserve Board of Governors.
- Bollerslev, T. and H. Zhou (2007), "Expected Stock Returns and Variance Risk Premia," Working Paper; Duke University and Federal Reserve Board of Governors.
- Bondarenko, O., 2000, "Recovering Risk-Neutral Densities: A New Nonparametric Approach," Working Paper; University of Illinois at Chicago.
- Bondarenko, O., 2003, "Estimation of Risk-Neutral Densities Using Positive Convolution Approximation," *Journal of Econometrics* 116, 85-112.
- Bondarenko, O., 2007, "Variance Trading and Market Price of Variance Risk," Working Paper; University of Illinois at Chicago.
- Breeden, D., and R. Litzenberger, 1978, "Prices of State Contingent Claims Implicit in Options Prices," *Journal of Business* 51, 621-652.
- Britten-Jones, M. and A. Neuberger, 2000, "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance* 55, 839-866.
- Carr, P., and D. Madan, 1998, "Towards a Theory of Volatility Trading," *Volatility*, Risk Publications, R. Jarrow, Ed., 417-427.
- Carr, P. and L. Wu, 2004, "Variance Risk Premia," Working Paper; New York University and Baruch College.
- Demeterfi, K., E. Derman, M. Kamal, and J. Zou, 1999, "More Than You Ever Wanted to Know About Volatility Swaps," <http://www.gs.com/qs/doc/volswaps.pdf>, Working Paper; Goldman Sachs.
- Duan, J.-C. and C.-Y. Yeh, 2007, "Jump and Volatility Risk Premiums Implied by VIX," Working Paper; Rotman School of Management, University of Toronto.
- Dumas, B., J. Fleming, and R. Whaley, 1998, "Implied Volatility Functions: Empirical Tests," *Journal of Finance* 53, 2059-2106.
- Dupire, B., 1993, "Model Art," *Risk* (September) 118-120.
- Jiang, G., and Y. Tian, 2005a, "The Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies* 18, 1305-1342.
- Jiang, G., and Y. Tian, 2005b, "Gauging the "Investor Fear Gauge": Implementation Problems in the CBOE's New Volatility Index and a Simple Solution," Working Paper; University of Arizona.
- Meddahi, N., 2002, "A Theoretical Comparison Between Integrated and Realized Volatility," *Journal of Applied Econometrics* 17, 479-508.

- Neuberger, A., 1994, "The Log Contract: A New Instrument to Hedge Volatility," *Journal of Portfolio Management*, Winter, 74-80.
- Ross, S., 1976, "Options and Efficiency," *Quarterly Journal of Economics* 90, 75-89.
- Todorov, V., 2007, "Variance Risk Premium Dynamics," Working Paper; Duke University.
- Whaley, R., 1986, "Valuation of American Futures Options: Theory and Empirical Tests," *Journal of Finance* 41, 127-150.
- Wu, L., 2004, "Variance Dynamics: Joint Evidence from Options and High-Frequency Returns," Working Paper; Baruch College.
- Zhang, L., P.A. Mykland and Y. Aït-Sahalia (2005), "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data," *Journal of the American Statistical Association* 100, 1394-1411.

Table 1: **Summary Statistics for Volatility**

Panel A: Full sample 01/1990-12/2006

	RVD	RVH	VIX	BSIV	CIV1	CIV2	CIV3	CIV4	MFIV
Mean	0.149	0.150	0.191	0.165	0.134	0.163	0.172	0.178	0.185
StDev	0.071	0.066	0.064	0.058	0.048	0.058	0.062	0.063	0.064
Skewness	1.500	1.428	0.983	1.017	1.009	1.042	1.064	1.069	1.080
Kurtosis	5.913	5.254	3.795	3.955	3.958	4.046	4.096	4.116	4.137
ρ_1	0.991	0.997	0.983	0.981	0.981	0.980	0.980	0.980	0.980
ρ_{21}	0.650	0.758	0.827	0.830	0.837	0.827	0.823	0.821	0.822
ρ_{63}	0.497	0.519	0.652	0.678	0.694	0.670	0.660	0.654	0.652

Panel B: Subsample 01/1990-12/1999

	RVD	RVH	VIX	BSIV	CIV1	CIV2	CIV3	CIV4	MFIV
Mean	0.138	0.139	0.185	0.157	0.126	0.155	0.164	0.170	0.177
StDev	0.063	0.057	0.059	0.051	0.041	0.052	0.055	0.057	0.058
Skewness	1.878	1.629	1.134	1.117	1.025	1.171	1.230	1.258	1.284
Kurtosis	8.983	7.132	4.526	4.830	4.590	5.048	5.182	5.277	5.327
ρ_1	0.988	0.996	0.979	0.974	0.974	0.973	0.974	0.975	0.975
ρ_{21}	0.586	0.736	0.806	0.799	0.805	0.796	0.792	0.790	0.794
ρ_{63}	0.413	0.467	0.635	0.651	0.667	0.643	0.633	0.626	0.627

Panel C: Subsample 01/2000-12/2006

	RVD	RVH	VIX	BSIV	CIV1	CIV2	CIV3	CIV4	MFIV
Mean	0.163	0.164	0.199	0.177	0.145	0.175	0.183	0.189	0.196
StDev	0.079	0.075	0.070	0.065	0.054	0.065	0.068	0.069	0.070
Skewness	1.091	1.107	0.760	0.775	0.774	0.793	0.804	0.799	0.803
Kurtosis	3.824	3.698	3.063	3.033	3.029	3.080	3.114	3.103	3.116
ρ_1	0.993	0.997	0.986	0.985	0.986	0.985	0.984	0.984	0.984
ρ_{21}	0.696	0.761	0.844	0.847	0.850	0.846	0.844	0.844	0.843
ρ_{63}	0.550	0.527	0.658	0.680	0.691	0.675	0.667	0.664	0.658

Notes: This table reports summary statistics for volatility measures RVD, RVH, VIX, BSIV, CIV1, CIV2, CIV3, CIV4, and MFIV. Statistics are reported for the full sample and two subsamples and include mean, standard deviation, skewness, kurtosis, and serial auto-correlation coefficients with lags of 1, 21, and 63 trading days. In all tables and figures, the volatility measures are annualized and given in decimal form.

Table 2: **Volatility Correlations**

	RVD	RVH	VIX	BSIV	CIV1	CIV2	CIV3	CIV4	MFIV
RVD	1.000	0.955	0.855	0.857	0.854	0.857	0.858	0.859	0.859
RVH	0.955	1.000	0.898	0.899	0.896	0.899	0.899	0.899	0.899
VIX	0.855	0.898	1.000	0.988	0.981	0.990	0.992	0.993	0.993
BSIV	0.857	0.899	0.988	1.000	0.998	1.000	0.998	0.997	0.995
CIV1	0.854	0.896	0.981	0.998	1.000	0.997	0.994	0.991	0.988
CIV2	0.857	0.899	0.990	1.000	0.997	1.000	0.999	0.998	0.996
CIV3	0.858	0.899	0.992	0.998	0.994	0.999	1.000	1.000	0.998
CIV4	0.859	0.899	0.993	0.997	0.991	0.998	1.000	1.000	0.999
MFIV	0.859	0.899	0.993	0.995	0.988	0.996	0.998	0.999	1.000

Notes: This table reports the correlations for various volatility measures RVD, RVH, VIX, BSIV, CIV1, CIV2, CIV3, CIV4, and MFIV. The sample period is 01/1990–12/2006.

Table 3: Volatility Regressions

	In-Sample Estimation				Out-of-Sample RMSE			
	α	β_1	β_2	R^2	All days	Low	Medium	High
RVD	0.05 (5.49)	0.62 (8.15)		46.75	31.91	16.47	26.81	44.15
RVH	0.04 (3.65)	0.74 (8.92)		54.05	30.23	13.91	25.62	42.12
VIX	0.00 (0.08)	0.75 (9.90)		60.25	27.98	10.52	26.45	38.08
BSIV	0.00 (0.36)	0.87 (10.79)		60.83	26.53	10.74	25.70	35.62
CIV1	0.00 (0.22)	1.09 (11.01)		60.79	26.11	10.78	25.39	34.97
CIV2	0.01 (0.55)	0.86 (10.79)		60.88	26.76	10.73	25.90	35.95
CIV3	0.01 (0.71)	0.80 (10.46)		60.59	27.15	10.72	26.21	36.56
CIV4	0.01 (0.70)	0.78 (10.34)		60.28	27.39	10.72	26.41	36.91
MFIV	0.01 (0.45)	0.76 (10.31)		60.34	27.58	10.83	26.58	37.18
RVD + RVH	0.04 (3.40)	-0.08 (-0.39)	0.82 (3.25)	54.13	30.28	13.88	25.72	42.17
RVH + VIX	0.01 (0.43)	0.19 (2.19)	0.59 (6.46)	60.91	27.42	10.53	25.70	37.40
RVH + CIV1	0.01 (0.60)	0.20 (2.04)	0.84 (8.26)	61.62	26.08	10.81	24.87	35.19
VIX + CIV1	0.00 (0.05)	0.30 (1.01)	0.67 (1.93)	61.15	26.47	10.51	25.54	35.64

Notes: The results refer to predictive and encompassing regressions for future monthly realized volatility, as measured by RVH. The explanatory (predictor) variables for each regression are listed in the left column. The estimation period is 01/1990–12/1999 and the out-of-sample forecast period is 01/2000–12/2006. Data are obtained for every trading day, so there is substantial overlap between successive observations on realized volatility. Heteroskedasticity and autocorrelation consistent t-statistics are reported below the regression coefficients.

Table 4: **Log-Volatility Regressions**

	In-Sample Estimation				Out-of-Sample RMSE			
	α	β_1	β_2	R^2	All days	Low	Medium	High
RVD	-0.69 (-6.77)	0.65 (13.97)		53.34	12.70	9.39	11.59	16.08
RVH	-0.44 (-4.09)	0.78 (15.62)		61.48	11.60	8.19	10.71	14.84
VIX	-0.31 (-2.79)	1.00 (16.39)		65.25	10.93	6.95	10.86	13.80
BSIV	-0.18 (-1.71)	0.98 (17.71)		66.79	10.34	6.82	10.56	12.66
CIV1	0.04 (0.30)	0.98 (17.93)		67.21	10.17	6.79	10.44	12.37
CIV2	-0.19 (-1.79)	0.97 (17.73)		66.80	10.42	6.84	10.65	12.79
CIV3	-0.26 (-2.48)	0.96 (17.39)		66.29	10.58	6.87	10.78	13.05
CIV4	-0.29 (-2.75)	0.96 (17.14)		65.81	10.68	6.91	10.88	13.20
MFIV	-0.29 (-2.76)	0.98 (17.10)		65.64	10.76	6.90	11.00	13.30
RVD + RVH	-0.44 (-4.02)	-0.02 (-0.21)	0.81 (6.14)	61.48	11.61	8.19	10.72	14.85
RVH + VIX	-0.28 (-2.59)	0.29 (4.18)	0.67 (7.18)	66.78	10.52	6.86	10.37	13.25
RVH + CIV1	-0.02 (-0.21)	0.23 (3.42)	0.73 (9.67)	68.13	10.07	6.79	10.17	12.35
VIX + CIV1	0.00 (0.03)	0.12 (0.49)	0.86 (4.12)	67.23	10.19	6.75	10.43	12.44

Notes: The results refer to predictive and encompassing regressions for future monthly log realized volatility, as measured by log RVH. The explanatory (predictor) variables for each regression are the log of the variable listed in the left column. The estimation period is 01/1990–12/1999 and the out-of-sample forecast period is 01/2000–12/2006. Data are obtained for every trading day, so there is substantial overlap between successive observations on realized log volatility. Heteroskedasticity and autocorrelation consistent t-statistics are reported below the regression coefficients.

Table 5: **Variance Regressions**

	In-Sample Estimation				Out-of-Sample RMSE			
	α	β_1	β_2	R^2	All days	Low	Medium	High
RVD	0.01 (5.12)	0.51 (5.14)		37.51	68.27	24.04	52.52	97.04
RVH	0.01 (3.70)	0.66 (5.31)		43.07	65.98	19.17	51.81	93.88
VIX	0.00 (0.08)	0.60 (5.66)		50.82	59.21	11.27	53.28	82.44
BSIV	0.00 (0.38)	0.80 (6.25)		50.18	56.95	12.00	52.14	78.89
CIV1	0.00 (0.27)	1.26 (6.35)		49.18	56.36	12.10	51.61	78.05
CIV2	0.00 (0.57)	0.80 (6.28)		50.45	57.39	12.07	52.50	79.52
CIV3	0.00 (0.69)	0.70 (6.08)		50.74	58.06	12.07	52.99	80.50
CIV4	0.00 (0.69)	0.65 (6.06)		50.81	58.42	12.03	53.25	81.04
MFIV	0.00 (0.50)	0.62 (6.10)		51.43	58.66	12.20	53.39	81.39
RVD + RVH	0.01 (3.26)	-0.01 (-0.04)	0.67 (1.64)	43.05	65.99	19.13	51.84	93.89
RVH + VIX	0.00 (0.32)	0.13 (1.21)	0.50 (4.44)	51.20	58.92	11.42	52.51	82.21
RVH + CIV1	0.00 (0.69)	0.23 (1.80)	0.92 (5.41)	50.63	57.08	12.35	50.80	79.57
VIX + CIV1	0.00 (0.06)	0.50 (1.50)	0.21 (0.39)	50.87	58.33	11.29	52.82	81.08

Notes: The results refer to predictive and encompassing regressions for future monthly realized return variation, as measured by the squared value of RVH. The explanatory (predictor) variables for each regression are the squared values of the variables listed in the left column. The estimation period is 01/1990–12/1999 and the out-of-sample forecast period is 01/2000–12/2006. Data are obtained for every trading day, so there is substantial overlap between successive observations on realized return variation. Heteroskedasticity and autocorrelation consistent t-statistics are reported below the regression coefficients.

Table 6: Volatility Regressions

	In-Sample Estimation				Out-of-Sample RMSE			
	α	β_1	β_2	R^2	All days	Low	Medium	High
RVD	0.05 (5.46)	0.71 (13.85)		54.70	26.00	30.04	21.83	19.83
RVH	0.04 (3.88)	0.77 (11.99)		57.94	23.86	27.99	19.23	18.24
VIX	-0.01 (-1.62)	0.89 (15.07)		67.96	23.46	27.05	19.73	17.98
BSIV	-0.01 (-0.91)	0.96 (15.19)		68.39	22.16	26.26	18.07	16.18
CIV1	-0.01 (-0.66)	1.16 (15.07)		68.52	21.93	26.21	17.61	15.87
CIV2	-0.01 (-0.74)	0.97 (15.16)		68.15	22.24	26.33	18.20	16.24
CIV3	-0.01 (-0.71)	0.92 (15.18)		67.93	22.57	26.66	18.56	16.47
CIV4	-0.01 (-0.79)	0.90 (15.14)		67.74	22.79	26.82	18.86	16.72
MFIV	-0.01 (-1.28)	0.89 (14.97)		67.42	22.82	26.79	18.92	16.85
RVD + RVH	0.04 (3.92)	0.09 (0.36)	0.68 (2.57)	57.98	23.94	28.09	19.37	18.23
RVH + VIX	-0.01 (-1.27)	0.11 (0.88)	0.79 (5.90)	68.18	23.05	26.78	19.30	17.35
RVH + CIV1	-0.00 (-0.45)	0.11 (0.81)	1.03 (5.34)	68.74	21.74	26.01	17.52	15.62
VIX + CIV1	-0.01 (-0.88)	0.23 (0.47)	0.87 (1.37)	68.58	22.00	26.10	17.86	16.08

Notes: The results refer to predictive and encompassing regressions for future monthly realized volatility, as measured by RVH. The explanatory (predictor) variables for each regression are listed in the left column. The estimation period is 01/2000–12/2006 and the out-of-sample forecast period is 01/1990–12/1999. Data are obtained for every trading day, so there is substantial overlap between successive observations on realized volatility. Heteroskedasticity and autocorrelation consistent t-statistics are reported below the regression coefficients.





Figure 1: The top panel plots the normalized barrier variance $U_0(p)$ for the Black-Scholes model, when $\sigma = 0.165$ and $\tau = 21$ trading days. The bottom panel plots the normalized barrier variance $U_0(p)$ for S&P 500 on 04/19/2000, when $\tau = 21$ trading days. On that day, S&P 500 = 1427.47, VIX = 27.02.

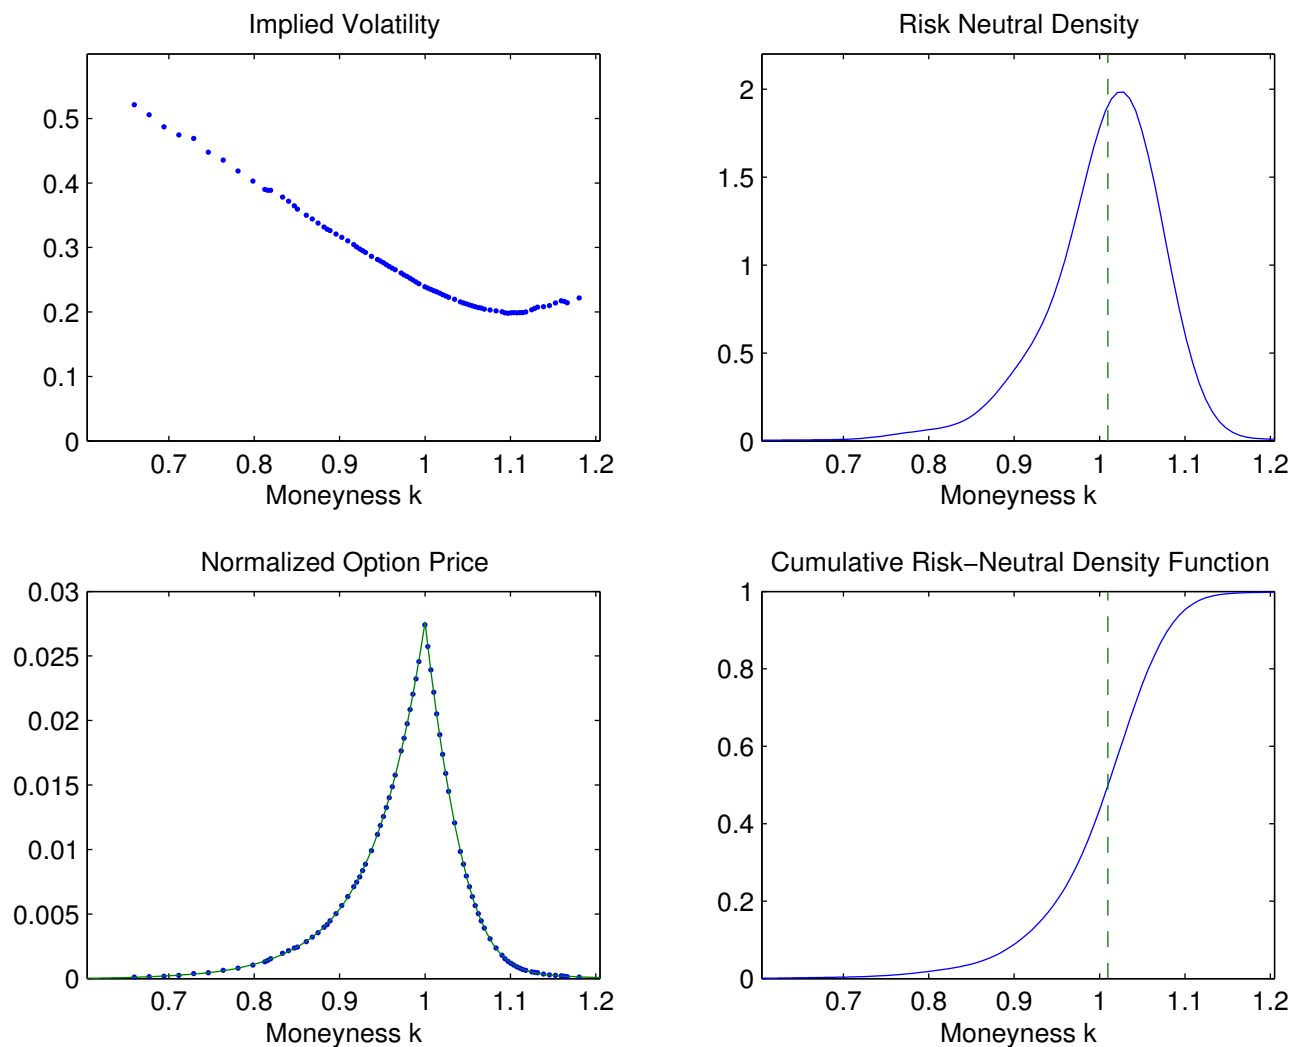


Figure 2: This figure illustrates S&P 500 option data for 04/19/2000 when $\tau = 21$ trading days. The left panels show the Black-Scholes implied volatility $IV_i(k)$ and the normalized OTM option price $M_t(k)/F_t = \min(P_t(k), C_t(k))/F_t$ for different values of moneyness $k = K/F_t$. The solid line indicates option prices corresponding to the estimated RND. The right panels show the estimated RND $h_t(k)$ and the cumulative risk-neutral density function $H_0(k)$. The dashed line indicates the median of the RND. On that day, S&P 500 = 1427.47, VIX = 27.02.




Figure 3: The top panel plots the level of S&P 500 index (log scale). The second panel plots daily returns on S&P 500 index. The bottom panel plots VIX (the thick line) and realized volatility RVH (the thin line). The realized volatility RVH is computed using high-frequency returns over 21 trading days. VIX and RVH are annualized and given in decimal form. The vertical dashed line separates the estimation and forecast periods.





Figure 4: The top panel plots the corridor variances CIV1-CIV4 scaled by MFIV. The bottom panel plots the corridor variances CIV1-CIV4 and MFIV scaled by BSIV.