Dissecting the Market Pricing of Return Volatility

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

May 2009

Abstract

The model-free implied volatility (MFVOL) has gained notoriety in recent years, partially due to the Chicago Board Option Exchange (CBOE) publication of the VIX index which is based on the MFVOL. This volatility measure extracts the expected future return volatility from the full cross-section of options with different strikes for a given maturity date. Conceptually, the MFVOL provides the market price of future volatility. If there is a volatility risk premium, the MFVOL will deviate systematically from the future realized return volatility. This is the case for the VIX, as it overstates the subsequent volatility by a large margin on average. This implies that exposure to S&P 500 volatility carries a negative risk premium. Intuitively, investors are willing to pay high prices for the insurance against spikes in the market volatility.

In this paper, we refine the analysis of the volatility risk premium in many dimensions. We exploit the concept of corridor implied volatility to break the MFVOL into slices representing the pricing of market volatility over different intervals for the underlying asset price. One, we study the size of the volatility risk premium for drops versus increases in the underlying asset price. Two, we extract the premium for volatility risk in the tails of the risk-neutral price density. Three, we explore asset classes beyond the broad equity market, in particular, foreign exchange markets. Four, we develop conditional measures of the time-variation in the relative pricing of risk in up- and down-states. This down-up-ratio (DUR) provides a gauge on the shifts in the pricing of volatility on the downside versus the upside. Our findings suggest that the DUR is highly correlated with the overall risk aversion in the market and thus provides a theoretically motivated summary statistic which speaks to the time-variation in the pricing of market risks more generally.

JEL Classification: C53, G12, G13

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

---

[*] 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

In recent years, markets for direct trading of realized return volatility have emerged. A critical theoretical backdrop for these developments is the *model-free (implied) volatility* (MFVOL) which can be computed from the full cross-section of out-of-the-money options written on the underlying asset price. This notion underlies the VIX index released by the Chicago Board of Option Exchange which, in theory, prices the S&P 500 equity return variation over the coming month under the martingale (or Q) measure. A parallel development in financial econometrics has established that we can obtain close approximations to the realized return variation under the actual (or P) measure via high-frequency data, see, e.g., Andersen, Bollerslev and Diebold (2007). This enables us to explore the joint behavior of the return variation under the two measures. Moreover, the market premium for variance risk, given by the difference of realized variation under Q and P, can be obtained. A variety of studies based on this approach has emerged, shedding preliminary light on the character and properties of the variance risk premium for the equity market.

Broadly speaking, there are three related areas of inquiry in the academic literature. First, there is a very large set of papers exploring the realized volatility (P measure) concept, its decomposition into continuous and jump parts, and the empirical use of these measures for return volatility modeling and forecasting, for asset and derivatives pricing, for portfolio selection, for risk management, for specification testing, and for estimation of the impact of news announcements. Second, there is growing interest in the time series of MFVOL per se, in particular as represented by the VIX. Since the VIX correlates strongly with a variety of credit spreads, with equity market downturns, and with business cycle variables in general, it has been labeled a "fear gauge." As such, it has become a benchmark for the degree of turmoil in the markets during the financial crises originating in 2007 but intensifying following the bankruptcy of Lehman Brothers in September 2008. Third, a smaller but rapidly evolving literature studies the volatility risk premium, i.e., the gap between MFVOL and (expected) realized volatility, along with its implications for asset pricing in a number of directions.

The current paper is related to the second and third categories in the above classification scheme. Our main purpose is to shed additional light on the nature of the variance risk premium. The MFVOL, or VIX, reflects both the expected actual future return variation (under the P measure) and the pricing of volatility risk, often rationalized by the strong correlation between volatility and negative market developments. We argue that, to the extent any component of the VIX, or MFVOL, is to be interpreted as a "fear index," it must first isolate the volatility risk premium from the rational expectation of actual future return variation. Second, it should focus on excess volatility risk premiums associated with negative economic ("scary") developments. For example, if the premium is symmetric in the sense that volatility risk is priced equally richly for the scenarios involving a rising equity market index as for those involving a declining index, then the root cause of the premium is not directly associated with any aggregate fear, but rather with the volatility risk per se. Hence, we need a formal procedure for separating the volatility risk premium into a premium for volatility risk associated with negative versus positive market developments. We propose a simple measure which focuses directly on the differential pricing of volatility across good and bad economic states. This is accomplished exploiting the notions of corridor and barrier variance. To the best of our knowledge, we provide the first empirical study of these measures for the broad asset classes investigated here.

The *corridor variance* decomposes the model-free implied return variation over a given horizon into the components of the expected (Q measure) return variation over specific price ranges for the underlying asset. Carr and Madan (1998) introduce the concept and demonstrate that it may be priced similarly to the regular variance swap contract. We focus on the market price of variance risk for the coming month on the downside versus upside, i.e., the expected return variation when the current futures price is below versus above the starting futures price of the month. Specifically, we denote the ratio of the price of variance on the downside relative to the price of variance on the upside the *Down-Up Ratio* or *DUR*. This is an intuitive, readily interpretable and theoretically coherent quantitative measure of the time-varying degree of asymmetry inherent in the skew of the Black-Scholes implied volatilities. While these computations are all performed under the Q measure, we implement the corresponding decomposition of actual realized return variation into the return variation incurred in down- and up-states (under the P measure) as well, utilizing high-frequency data on the underlying asset. With these tools in hand, we are able to investigate not only the overall size of the risk premium, but also where the risk premium resides within the risk neutral density and, in particular, we obtain separate volatility risk premiums for the down-state and the up-state.

The determination of a market price (Q measure expectation) of the amount of return variation incurred during periods when the market is below the level at the start of the month is related to the so-called semi-variance. This is a variance measure computed using return observations falling below a given threshold only. Semi-variance is arguably the most relevant variance-type measure for downside risk as it speaks to the size of average losses below a threshold rather than simply the average size of price moves. The corresponding downside realized volatility measure, computed from high-frequency return data, provides a close approximation to the actual return variation incurred in the down region. It may serves as a measure of the ex-post realization corresponding to the ex-ante semi-variance expectation. In this sense, it parallels the theoretical relation between the ex-post realized volatility measure and the *a priori* expected return variance, except that the focus here, exclusively, is on the downside (or upside) return variation.

Since we construct the various MFVOL measures from first principles rather than relying on a published VIX type index, we may extend the analysis to other asset classes of interest and study the premiums across separate maturities as long as data of sufficient quality are available. We focus our auxiliary analysis on the foreign exchange market, studying the pricing of volatility risk inherent in long time series of U.S. dollar futures and futures options prices for yen and British pound.

The empirical findings are quite striking. The mean down-up ratio, or DUR, for the sample is over 2.31, so the implied option market pricing of realized variance on the downside exceeds that on the upside by more than 136 per cent. In contrast, the actual realized return variation in "down states" exceeds that of "up states" by about 74 per cent so the variance risk premium is large and negative. Expressed in terms of a monthly percentage risk premium, it averages close to 33 per cent over the sample, but the premium is much larger (more negative) for downside than upside volatility, attaining values of 41 and 13 per cent respectively over the 21 year long sample. The main new finding is not the existence of a large variance risk premium but the quantification of the premium on the downside versus upside in a readily interpretable and economically meaningful metric. These findings suggest that downside variance risk, or semi-variance, incorporates, or loads heavily on, a risk factor which is very richly compensated in the options markets.

For the foreign exchange markets, we observe features that resemble the qualitative patterns in the equity market, although the (negative) volatility risk premiums are smaller and less asymmetric. Also in this case, we find that the realized and model-free volatilities are skewed in that, overall, both currencies have been trending in one direction while the realized volatility is dominated by relatively large eruptions in the opposite direction. This parallels the behavior of the equity index which has trended upward but displayed (priced) excess return variation to the downside. Finally, we note that, for both currencies, the volatility risk premiums remain negative on both the upside and downside, with an overall size of 14 per cent for the yen and 17 per cent for the pound.

Turning to the time series behavior of the DUR statistic, we identify sharp shifts around dramatic market events. The statistic for the equity index jumps upward around episodes such as the mini stock market crash in 1989, the Asian currency crises in 1997, the Russia credit event, the LTCM demise in 1998, and the default of Lehman Brothers in 2008. In contrast, it drops significantly upon the successful launch of the U.S. war effort in Kuwait in 1991 which arguably reduced overall volatility as well as "fear." We also provide a brief characterization of the dramatic rise in the volatility measures experienced during the latter part of 2008. In short, the overall level of volatility attained values far beyond what has been observed previously while "fear", as represented by the proportional pricing of downside to upside variance risk, merely hit a historical high.

Finally, we note that the DUR statistic displays persistent shifts across the sample. The equity-index DUR series appears stationary but it is common for the series to remain above or below the long term average of 2.31 for a year or more. In contrast, the currencies appear to experience a sharp change, reminiscent of a structural break, in their DUR during 2007. The yen DUR shifts sharply downward while the pound DUR shows a dramatic increase. We argue this is associated with the breakdown of the carry trade and the unwinding of short yen and long pound positions. This is also a period in which the global capital markets often are characterized as experiencing a sharp drop in the aggregate risk tolerance.

Finally, we also find that the volatility risk premiums are particularly rich for the more extreme tails of the risk-neutral distribution although this is more evident for the equity index than the currencies.

The remainder of the paper evolves as follows. Section 2 presents the theoretical backdrop to the analysis, including an account of the model-free barrier and corridor variance concepts and the use of corresponding realized volatility measures. Section 3 details the data exploited in the empirical analysis. Section 4 presents the results, and Section 5 provides some 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 to convert the 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 underlying 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 exposited in Ross (1976), Breeden and Litzenberger (1978), and Banz and Miller (1978),

$$h_t(F_T) = \left. \frac{\partial^2 P_t(K)}{\partial K^2} \right|_{K=F_T} = \left. \frac{\partial^2 C_t(K)}{\partial K^2} \right|_{K=F_T} \tag{1}$$

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:

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

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

$$\begin{aligned} E_0^Q \left[g(F_T) \right] &= g(F_0) + \int_0^{F_0} g''(K) P_0(K) dK + \int_{F_0}^\infty g''(K) C_0(K) dK \\ &= g(F_0) + \int_0^\infty g''(K) M_0(K) dK, \end{aligned} \tag{3}$$

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.

---

[1] In 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).

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

$$\frac{dF_t}{F_t} = \sigma_t dW_t, \tag{4}$$

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,

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

which implies that

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

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

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

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*,

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

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:

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

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

$$MFBV_0(B) = E_0^Q\left[\int_0^T \sigma_t^2 I_t dt\right] = 2\int_0^B \frac{M_0(K)}{K^2} dK, \tag{8}$$

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

$$MFBVOL_0(B) = \sqrt{2\int_0^B \frac{M_0(K)}{K^2} dK}, \tag{9}$$

In the limiting case of $B = \infty$, the barrier implied volatility coincides with the model-free implied volatility $MFVOL_0$. 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,

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

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

$$MFCV_0(B_1, B_2) = E_0^Q\left[\int_0^T \sigma_t^2 I_t dt\right] = 2\int_{B_1}^{B_2} \frac{M_0(K)}{K^2} dK. \tag{10}$$

2.2 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

---

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

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 $0 = t(0) < t(1) < \ldots < t(n-1) < t(n) = T$, we may define realized volatility as the cumulative sum of squared returns,

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

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.,

$$V_0 = E_0^P \left[\int_0^T \sigma_t^2 dt \right] = E_0^P \left[RV_n(0,T) \right] . \tag{11}$$

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 semi-martingale 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. As a result one may prefer to use a slightly lower frequency, such as evenly spaced 5-minute returns, to construct the RV measure. Given the limited microstructure effects in these futures markets, the analysis of Andersen, Bollerslev and Meddahi (2006) indicates that this should work quite well in practice. Nonetheless, this will not necessarily eliminate the full bias and it inevitably entails a loss of efficiency. Consequently, we follow the so-called sub-sampling procedure of Zhang, Mykland and Aït-Sahalia (2005)which aggregates lower frequency squared futures returns across the trading day, but does so repeatedly using a different starting point for the trading day so that distinct overlapping returns are used for each separate computation. Hence, we obtain multiple RV measures for the given trading day which then are averaged to produce the final sub-sampled RV. This approach allows for a more efficient use of the data while operating at a sufficiently low frequency that we obtain an approximately unbiased measure.[4] Specifically, we rely on 15-minute returns for the S&P 500 futures market and 30-minute returns for the currency futures markets using 15 subsamples in each case. Hence, we effectively exploit 1-minute equity index and 2-minute foreign exchange returns.

Finally, we also need to account for 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 sub-sampled measure obtained over the trading period. 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.

---

[3] This 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).

[4] For a formal account of this approach and related procedures, see Zhang, Mykland and Aït-Sahalia (2005), Bandi and Russell (2005), and Barndorff-Nielsen, Hansen, Lunde and Shephard (2008).

2.3 Construction of the Variance Measures

In the empirical study, we compare and contrast properties of the various 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 construct model-free Down and Up variance, denoted MFVD and MFVU as,

$$
\begin{aligned}
MFVD_0 &= 2\int_0^{F_0} \frac{M_0(K)}{K^2} dK, \\
MFVU_0 &= 2\int_{F_0}^{\infty} \frac{M_0(K)}{K^2} dK.
\end{aligned}
$$

The model-free (total) variance, MFV, is simply

$$MFV_0 = MFVD_0 + MFVU_0.$$

To construct the barrier model-free variances (MFVD and MFVU), 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.

3 Data

The full sample period is from January 1988 through December 2008. From the Chicago Mercantile Exchange (CME), we obtain transactions data for the futures on three assets (S&P 500 index, Japanese yen, and British pound) as well as daily prices for the corresponding options. From the U.S. Federal Reserve, we obtain Treasury bill rates, which are used to proxy for the risk-free interest rate.

3.1 S&P 500

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, a 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.

3.2 Japanese Yen and Britich Pound

Similarly to the S&P 500 futures, the Japanese Yen (JY) and British Pound (BP) futures also have four different maturity months from the March quarterly cycle. On any trading day, the futures options are available for six maturity months: four months from the March quarterly cycle and two additional nearby months ("serial" options).

Both the JY and BP options expire on the second Friday preceding the third Wednesday of a contract month. The JY futures contract size is 12,500,000 Japanese yen. The JY option contract size is one JY futures. The minimum price movement is $.000001 per Japanese yen increments (or $12.50/contract). The BP futures contract size is 62,500 British pounds. The BP option contract size is one BP futures. The minimum price movement is $.0001 per British Pound increments (or $6.25/contract).

When computing RV measures for JY and BP, we use transactions data for the futures from both pit and electronic trading. The two trading venues have co-existed over a long period in our sample, however, around 2002 to 2005 liquidity in FX market steadily migrated from pit trading to electronic one. For that reason, we use pit data from 1988 to 2004 and electronic trading data after that. The data source for each futures series is summarized in Table 1.

For each trading day, we estimate the model-free variances, MFVD, MFVU, and MFV. To obtain these values, we follow several steps, which are described in more detail in Appendix A. 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 1 depicts the Black-Scholes implied volatility, the normalized option prices, and the risk-neutral pdf and cdf for a representative trading day in the sample.

The corresponding realized variance measures, RVD, RVU, and RV are computed from futures transactions data using the subsampling procedure as explained in detail previously.

4 Empirical Results

We provide some simple summary statistics for our series in Tables 2 and 3. The first column of Table 2 reports the average fraction of time over the ensuing month the futures price spends either above or below the futures price at the beginning of the month. Since a constant futures price for the SP500 over the life of the holding period is consistent with increases in the underlying spot index at the risk-free rate, the positive tilt in the distribution reflects the mean excess return of the equity index over the risk-free rate. For the currency futures, the reference growth in the underlying spot exchange rate, consistent with no change in the futures price, is given by the (risk-free) interest rate differential between the currencies. Hence, the predominance of downward moves for the yen and upward moves for the pound implies that the carry trade of shorting yen and going long pound (both relative to the opposite position in the U.S. dollar) would have been profitable most of the time across the sample period. It does not imply that the yen actually depreciated over the sample - it did not - but rather that it typically failed to appreciate by as much as indicated by the interest rate differential.

The remaining columns in Table 2 provide sample moments for the realized and model-free variance measures. We discuss some of these in more detail below, but we want to draw attention to the apparent and interesting asymmetries for the up- and down-variances across all these assets. In each case, the realized variance measure is highest for the state in which the futures price spend the least amount of time - for the equity index and the pound, the realized down-variance is substantially larger than the up-variance, while for the yen the up-variance exceeds the down-variance. This suggests that the price process is relatively tranquil in the more common state and more volatile in the less usual state. This is also evident in the corresponding standard deviation measures for the up- and down-states. Furthermore, this tendency to realize more return variation in the unusual state is also priced into the options, as the model-free up- and down-variances display the same ordering for all assets.

Table 3 reports correlations across the various measures. It is evident that the model-free implied variance measures all have predictive power for subsequent realized return variation. This is consistent with the general findings that implied volatility measures do carry important information regarding future return volatility.[5] Moreover, the (total) realized variance is, almost by construction, highly correlated with the up-variance and down-variance measures. Finally, as expected, the correlation between the realized up-variance and down-variance measures is close to zero or negative. While both measures will be high, on average, whenever the overall realized variance is high, the realized up-variance also will tend to be higher the more time the futures price spends in the up-state rather than the down-state over the course of the month.

Tables 4-8 provide additional evidence on the variance measures and the volatility risk premium as well as on the composition of the premium across different economic states. The empirical findings are striking. Combining Tables 2 and 4, we find that for the benchmark one-month horizon in the U.S. equity-index futures market, the overall down-up ratio, or DUR,

---

[5] Andersen and Bondarenko (2007) analyze the relative information content of many different model-free corridor implied variance measures, finding important dispersion in performance across these measures.

computed over the full sample period, 1988-2008, is in excess of 2.36 while the average monthly DUR is 2.31.[6] Hence, the implied option market pricing of realized variance on the downside exceeds that on the upside by more than 136 per cent over the full sample and by about 131 per cent on average across the months. In contrast, over the sample, the actual realized return variation experienced in "down states" exceeds that of "up states" by about 74 percent. Hence, while the actual realized return variation is highly negatively skewed, the model-free measures imply an even more dramatic skew in the volatility distribution. Expressed within a slightly different metric, Table 4 shows that the standard (overall) variance risk premium is large and negative at close to 33 per cent over the sample, but the premium is much larger (more negative) for downside than upside volatility, attaining values of 41 and 13 per cent respectively over the 21 year long sample.

This provides a couple of initial insights. First, there is a very substantial return-volatility asymmetry, or leverage effect, in the equity-index market with the return variation on the downside exceeding that on the upside by a huge margin. This implies that there should be a sizeable gap between the pricing of return variation on the downside versus upside. However, second, the option implied variance measures incorporate even more extreme premiums, implying a very large variance risk premium (in excess of the actual expected return variation under the P measure) for down-states relative to up-states. The novelty here is not the existence of a large variance risk premium but the quantification of the premium on the downside versus upside in an economically interpretable metric. These findings are obtained from sub-sampled realized volatility measures and carefully calibrated MFVOL measures based on a longer sample than hitherto analyzed. These results suggest that downside variance risk, or semi-variance, incorporates, or loads heavily on, a risk factor which is richly compensated in the options markets. Figure 3 plots the time series realizations of the model-free and realized equity-index variances across the sample, illustrating the size and variability of the variance risk premium.

For the foreign exchange markets, we observe features that resemble the qualitative patterns in the equity market. As previously observed, for the yen the majority of the realized volatility is incurred in the upstate. Consistent with this observation, the mean DUR value for the sample is below unity so that on average volatility on the upside has been priced more richly than volatility on the downside. The British pound provides a mirror image of the yen with both higher realized and model-free volatility being incurred on the downside so that the mean DUR value is well in excess of unity. Hence, overall, both currencies have been trending in one direction while the realized volatility is dominated by relatively large eruptions in the opposite direction. Finally, we note that, for both currencies, the volatility risk premiums remain negative on both the upside and downside, with an overall mean monthly premium on the order of -9 per cent for yen and -15 per cent for pound. The time series of model-free and realized variances for the currencies are plotted on Figures 7 and 11.

Turning to the time series behavior of the DUR statistic, we identify sharp shifts around dramatic market events in Figures 2, 6 and 10. The statistic for the equity index jumps upward around episodes such as the mini stock market crash in 1989, the Asian currency crises in 1997, the Russia credit event and the LTCM demise in 1998, and the default of Lehman Brothers in 2008. In contrast, it drops significantly upon the successful launch of the U.S. war effort in Kuwait in 1991 which arguably reduced overall volatility as well as "fear."

---

[6] The ratio constructed from the full sample values will be larger than the mean monthly ratio as the DUR is a concave function of the RV and MFVOL measures.

The evidence also provides a more detailed characterization of the dramatic rise in the volatility measures experienced during the latter part of 2008. The DUR statistic jumps very substantially but the ratio is not out of line with observations during previous crisis such as the LTCM episode in 1998. Hence, what truly distinguishes recent events is the overall level of volatility which attains values far beyond what has been observed previously. In other words, the dramatically elevated MFVOL and VIX readings reflect the fact that market uncertainty shot up to an unprecedented level while "fear", as represented by the proportional pricing of downside to upside variance risk, merely hit a historical high but not a plateau signifying a dramatic departure from earlier experiences. Figure 5 depicts the relevant development in these series following the Lehman Brothers default.

The time series behavior of the DUR statistic for the foreign exchange markets behave roughly similarly to the equity-index DUR series. It jumps sharply for the yen during the Asian currency crises, reflecting expectations, and perhaps fear, of higher downside yen volatility. The aftermath of September 11, 2001, provides an illustration of increasing volatility expectations, as the overall MFVOL for yen jumps very significantly, but only a minor "fear" shift as the yen DUR measure drops only marginally. For the British pound, the departure from the EMS system in 1992 was accompanied by a jump in the DUR and higher overall volatility. An even larger, albeit also more gradual, reaction occurred during the financial crisis the Fall of 2008. A couple of these episodes are highlighted in Figures 9 and 13.

In order to gauge whether the downside variance risk premium has a "fear" component, it is useful to consider the decomposition of the variance pricing into states where the down-variance is priced higher than usual, relative to the up-variance, and when the opposite is true. That is, we study the realized variance and the associated variance risk premium across states where the DUR is above versus below its average value. If fear (or sharply time-varying premiums for downside variance risk) is an important factor, the (ex post) variance risk premium should be significantly larger (more negative) when DUR is high than when it is low. Table 5 reveals that this, indeed, appears to be the case for the equity index as the realized down- and up-variances are statistically indistinguishable in the high versus low DUR state. In other words, there is no evidence of elevated down-variance in states where the down-variance is highly priced relative to the up-variance. Also, note that the high equity DUR values reflect higher MFVD figures rather than lower MFVU figures as the mean MFVU is identical across those two groupings. This suggests that high equity DUR states largely reflect variation in the relative premium for down-variance and not any actual increase in the return variation on the downside versus the upside.

In contrast, the corresponding evidence for the currencies in Table 5 does not point towards any excessive risk premiums during periods where the DUR takes on extreme values, as the relative premiums for upside versus downside return variation across the low and high DUR states are quite similar and display no statistical differences.

Finally, we notice that the DUR statistic displays fairly pronounced and persistent shifts across the sample. The equity-index DUR series appears stationary but it is common for the series to remain above or below the long term average of 2.31 for a year or more. Interestingly, the currencies provide a marked contrast as there is a shift, reminiscent of a structural break, in their DUR series at some point during 2007. The yen DUR shifts sharply downward while the pound DUR shows a dramatic increase, in both cases to levels never previously observed in the sample. This is likely associated with the breakdown of the carry trade and the unwinding of short yen and long pound positions. This is also a period in which the global capital markets often are characterized as experiencing a sharp drop in the aggregate risk tolerance. In that context, the rapid ascent of the price of hedges against yen upside volatility and pound downside volatility is consistent with a dramatic increase in the expectation, and fear, of a yen appreciation and pound depreciation, respectively. As such, the currency DUR statistics provide a direct window into the current perception of these forces in the capital markets.

The paper provides some additional refinements of the analysis. For example, from Table 8, we note that the volatility risk premiums are particularly rich for the tails of the risk-neutral distribution (the regions V1 and V4 correspond to the lower and upper 25 per cent quartiles of the RND for the futures price one-month ahead). Hence, the premiums load heavily, although far from exclusively, on the tail events. This is particularly evident for the extreme down volatility in the equity index, while it is less striking for the currencies where the premiums appear closer to uniformly distributed across the domain of the underlying currency values. Nonetheless, the latter also appear to have mildly elevated volatility premiums in the tails. The relative sizes of these quartiles are also captured graphically on the bottom panels of Figures 4, 8, and 12.

5 Conclusion

This paper provides the first empirical study of the risk premiums for upside and downside variance across three major assets. We provide a variety of empirical findings that should stimulate future work in this area.

Appendix

A 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]

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

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

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.

---

[7] As 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.

[8] It 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.

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.

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. and O. Bondarenko, 2007, "Construction and Interpretation of Model-Free Implied Volatility," in Israel Nelken (Ed.): Volatility as an Asset Class, pp. 141-181; Risk Books, London.

Andersen, T.G., P.H. Frederiksen and A. Staal, 2006, "The Information Content of Realized Volatility Forecass," 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, 2008, "Designing Realised Kernels to Measure the Ex-Post Variation of Equity Prices in the Presence of Noise," *Econometrica* 76, 1481-1536.

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: **Data Description**

	Abbr.	Source	Sample Period Pit	Sample Period Electronic
S&P 500	SP	CME	01/1988-12/2008	n/a
Japanese Yen	JY	CME	01/1988-12/2003	01/2004-12/2008
British Pound	BP	CME	01/1991-12/2003	01/2004-12/2008

Table 2: **Summary Statistics for Realized and Model-Free Variances**

	% of returns	Realized Variance × 100 Mean	Realized Variance × 100 StDev	Realized Variance × 100 Skew.	Realized Variance × 100 Kurt.	Model-Free Variance × 100 Mean	Model-Free Variance × 100 StDev	Model-Free Variance × 100 Skew.	Model-Free Variance × 100 Kurt.
S&P 500									
Down	44.19	1.97	4.71	7.89	85.15	3.12	3.55	6.83	74.32
Up	55.81	1.13	1.68	7.87	113.53	1.32	1.22	5.05	46.19
Total	100.00	3.10	5.19	7.24	70.45	4.44	4.74	6.37	66.46
Japanese Yen									
Down	53.86	0.58	0.59	2.29	13.09	0.72	0.46	3.55	31.43
Up	46.14	0.71	1.14	4.93	36.51	0.75	0.53	5.02	46.57
Total	100.00	1.29	1.20	4.62	32.90	1.47	0.97	4.15	35.99
British Pound									
Down	47.07	0.47	0.84	6.01	50.37	0.55	0.49	6.02	53.15
Up	52.93	0.40	0.43	3.40	24.80	0.49	0.35	4.23	30.80
Total	100.00	0.87	0.92	5.37	40.04	1.03	0.84	5.25	43.20

Notes: This table reports summary statistics for monthly measures of realized variances (RVD, RVU, and RVT) and model-free variances (MFVD, MFVU, and MFVT) for the three assets (S&P 500, Japanese Yen, and British Pound). The second column reports the percent of all returns used to compute the corresponding variance measure.

Table 3: **Correlations for Realized and Model-Free Variances**

S&P 500

	RVD	RVU	RV	MFVD	MFVU	MFV
RVD	1.00					
RVU	0.12	1.00				
RVT	0.95	0.43	1.00			
MFVD	0.49	0.72	0.68	1.00		
MFVU	0.51	0.74	0.71	0.97	1.00	
MFVT	0.50	0.73	0.69	1.00	0.98	1.00

Japanese Yen

	RVD	RVU	RV	MFVD	MFVU	MFV
RVD	1.00					
RVU	-0.16	1.00				
RVT	0.34	0.87	1.00			
MFVD	0.43	0.39	0.58	1.00		
MFVU	0.36	0.43	0.59	0.90	1.00	
MFVT	0.40	0.42	0.60	0.97	0.98	1.00

British Pound

	RVD	RVU	RV	MFVD	MFVU	MFV
RVD	1.00					
RVU	-0.06	1.00				
RVT	0.89	0.41	1.00			
MFVD	0.61	0.48	0.78	1.00		
MFVU	0.60	0.47	0.77	0.98	1.00	
MFVT	0.61	0.48	0.78	1.00	0.99	1.00

Notes: This table reports correlations of monthly measures of realized variances (RVD, RVU, and RV) and model-free variances (MFVD, MFVU, and MFV) for the three assets (S&P 500, Japanese Yen, and British Pound).

Table 4: **Risk Premium**

	Mean	$(RV\text{-}MFV) \times 100$			$\left(\frac{RV}{MFV}-1\right) \times 100$		
	DUR	Down	Up	Total	Down	Up	Total
				S&P 500			
All	2.31	-1.15	-0.19	-1.34	-40.96	-13.18	-32.77
		0.21	0.04	0.21	4.28	2.76	2.71
Low DUR	2.01	-0.59	-0.21	-0.80	-29.44	-15.73	-24.70
		0.29	0.06	0.28	7.03	3.85	4.43
High DUR	2.62	-1.71	-0.17	-1.88	-52.48	-10.64	-40.85
		0.24	0.05	0.24	3.85	3.61	2.44
				Japanese Yen			
All	0.97	-0.14	-0.04	-0.17	-13.05	-5.19	-8.86
		0.03	0.05	0.05	3.89	5.32	2.67
Low DUR	0.84	-0.11	-0.05	-0.16	-11.12	-7.13	-8.56
		0.03	0.06	0.06	4.31	6.62	3.22
High DUR	1.10	-0.16	-0.02	-0.19	-14.99	-3.25	-9.16
		0.04	0.08	0.08	6.24	7.77	4.08
				British Pound			
All	1.10	-0.08	-0.08	-0.16	-19.93	-9.50	-14.84
		0.04	0.02	0.03	4.07	4.04	2.22
Low DUR	1.04	-0.07	-0.05	-0.12	-18.49	-5.19	-12.03
		0.02	0.02	0.02	4.39	6.29	2.58
High DUR	1.16	-0.08	-0.12	-0.20	-21.37	-13.82	-17.65
		0.07	0.03	0.06	5.89	3.63	3.10

Notes: This table documents risk premiums for monthly down, up, and total variances of the three assets (S&P 500, Japanese Yen, and British Pound). Two specifications are considered: the constant and relative risk premiums. All trading days in the sample period are split into two halves corresponding to low and high levels of DUR, where $DUR = MFVD/MFVU$ is the model-free down-to-up variance ratio. The risk premiums are reported separately for (1) all days, (2) days with low DUR, and (3) days with high DUR. For each of the three cases, the second row contains HAC-corrected standard deviations. The second column reports the average DUR.

Table 5: **Realized and Model-Free Variances for Different DUR**

	$RV \times 100$			$MFV \times 100$		
	Down	Up	Total	Down	Up	Total
			S&P 500			
All	1.97	1.13	3.10	3.12	1.32	4.44
	0.28	0.07	0.32	0.21	0.07	0.28
Low DUR	2.06	1.11	3.17	2.65	1.32	3.97
	0.32	0.09	0.35	0.15	0.07	0.23
High DUR	1.88	1.15	3.03	3.59	1.32	4.91
	0.38	0.12	0.47	0.37	0.12	0.50
			Japanese Yen			
All	0.58	0.71	1.29	0.72	0.75	1.47
	0.03	0.06	0.07	0.03	0.03	0.06
Low DUR	0.50	0.71	1.21	0.61	0.76	1.37
	0.03	0.09	0.09	0.03	0.05	0.08
High DUR	0.66	0.72	1.38	0.83	0.74	1.57
	0.05	0.09	0.11	0.04	0.04	0.08
			British Pound			
All	0.47	0.40	0.87	0.55	0.49	1.03
	0.05	0.02	0.06	0.03	0.02	0.06
Low DUR	0.34	0.35	0.69	0.41	0.40	0.81
	0.02	0.02	0.03	0.02	0.02	0.03
High DUR	0.60	0.45	1.05	0.68	0.57	1.25
	0.10	0.03	0.11	0.06	0.04	0.10

Notes: This table reports mean realized variances (RVD, RVU, and RVT) and model-free variances (MFVD, MFVU, and MFVT) for the three assets (S&P 500, Japanese Yen, and British Pound). All trading days in the sample period are split into two halves corresponding to low and high levels of DUR, where $DUR = MFVD/MFVU$ is the model-free down-to-up variance ratio. Statistics are reported separately for (1) all days, (2) days with low DUR, and (3) days with high DUR. For each of the three cases, the second row contains HAC-corrected standard deviations.

Table 6: **Distribution of Monthly Returns for Underlying, Down Variance, Up Variance, and Total Variance**

	Min.	1%	5%	10%	Med.	90%	95%	99%	Max.
					S&P 500				
r_f	-31.25	-12.48	-7.11	-4.82	0.75	5.20	6.78	10.58	19.60
r_{vd}	-99.99	-99.86	-99.38	-98.36	-64.09	34.28	98.82	288.39	899.44
r_{vu}	-99.98	-99.62	-97.27	-91.90	-14.41	61.79	87.21	160.26	399.39
r_{vt}	-85.46	-77.94	-69.45	-65.75	-45.17	7.64	49.79	171.52	594.51
					Japanese Yen				
r_f	-4.23	-1.71	-1.07	-0.81	-0.02	0.79	1.15	2.03	8.61
r_{vd}	-99.99	-99.65	-98.11	-95.05	-26.75	76.58	119.60	257.86	867.17
r_{vu}	-100.00	-99.82	-98.90	-96.85	-37.67	125.43	194.86	409.76	614.55
r_{vt}	-82.53	-69.83	-60.92	-53.60	-22.07	47.63	88.61	196.03	453.83
					British Pound				
r_f	-4.40	-1.71	-0.95	-0.69	0.00	0.70	0.95	1.53	3.52
r_{vd}	-99.97	-99.80	-98.91	-96.88	-38.37	71.70	120.09	266.97	634.37
r_{vu}	-100.00	-99.67	-98.04	-94.19	-21.60	80.73	120.38	287.87	636.57
r_{vt}	-74.87	-65.77	-55.69	-49.67	-22.71	21.45	51.18	160.26	299.84

Notes: Statistics are reported for excess returns of the underlying futures, Down variance contract, Up variance contract, and Total variance contract for the three assets (S&P 500, Japanese Yen, and British Pound). Returns are expressed in monthly percentage term. Statistics include minimum, median, maximum, and 1%, 5%, 10%, 90%, 95%, and 99% percentiles.

Table 7: **Risk Characteristics of Monthly Returns for Underlying, Down Variance, Up Variance, and Total Variance**

	Mean	SD	Skew.	Kurt.	α	β	SR
				S&P 500			
r_f	0.38	4.41	-0.75	6.69	0.00	1.00	0.09
r_{vd}	-40.96	80.70	3.80	25.82	-36.19	-12.59	-0.51
r_{vu}	-13.18	62.22	0.99	5.88	-16.19	7.95	-0.21
r_{vt}	-32.77	49.32	4.41	32.83	-30.42	-6.22	-0.66
				Japanese Yen			
r_f	-0.00	0.72	0.67	10.59	0.00	1.00	-0.01
r_{vd}	-13.05	80.42	2.50	18.20	-13.06	-1.95	-0.16
r_{vu}	-5.19	105.07	1.98	8.23	-5.17	5.17	-0.05
r_{vt}	-8.86	51.01	2.35	11.42	-8.85	1.84	-0.17
				British Pound			
r_f	0.00	0.60	-0.16	6.53	0.00	1.00	0.00
r_{vd}	-19.93	79.28	1.94	9.79	-19.92	-3.17	-0.25
r_{vu}	-9.50	80.60	1.93	10.49	-9.51	1.64	-0.12
r_{vt}	-14.84	38.69	2.61	13.74	-14.84	-0.97	-0.38

Notes: Statistics are reported for excess returns of the underlying futures, Down variance contract, Up variance contract, and Total variance contract for the three assets (S&P 500, Japanese Yen, and British Pound). Returns are expressed in monthly percentage term. Statistics include mean, standard deviation (SD), skewness, kurtosis, α and β coefficients (with respect to the underlying futures), Sharpe ratio (SR).

Table 8: **Summary Statistics for Realized and Model-Free Variances V1-V4**

	% of returns	Realized Variance × 100				Model-Free Variance × 100			
		Mean	StDev	Skew.	Kurt.	Mean	StDev	Skew.	Kurt.
S&P 500									
V1	15.05	1.09	3.88	9.75	127.79	2.03	2.43	7.22	81.73
V2	29.14	0.88	1.61	9.03	123.60	1.08	1.14	6.03	60.90
V3	36.92	0.79	1.27	9.70	160.13	0.90	0.88	5.45	51.96
V4	18.89	0.34	0.76	6.63	89.08	0.42	0.35	4.02	32.07
Japanese Yen									
V1	22.50	0.24	0.40	2.64	11.94	0.34	0.23	3.34	29.82
V2	31.36	0.34	0.36	3.89	37.44	0.37	0.25	4.11	36.69
V3	26.88	0.35	0.44	4.35	33.59	0.37	0.25	4.33	37.80
V4	19.26	0.37	0.91	6.18	55.61	0.38	0.29	5.60	54.95
British Pound									
V1	18.37	0.23	0.69	7.36	72.55	0.27	0.27	6.80	63.74
V2	28.70	0.24	0.30	4.95	43.28	0.28	0.23	5.13	41.92
V3	31.46	0.24	0.28	5.57	57.32	0.27	0.21	4.73	36.76
V4	21.47	0.17	0.27	2.43	10.45	0.22	0.14	3.51	22.83

Notes: This table reports summary statistics for monthly measures of realized variances (RV1, RV2, RV3, and RV4) and model-free variances (MFV1, MFV2, MFV3, and MFV4) for the three assets (S&P 500, Japanese Yen, and British Pound). The second column reports the percent of all returns used to compute the corresponding variance measure.

Table 9: **Correlations for Realized and Model-Free Variances V1-V4**

S&P 500

	RV1	RV2	RV3	RV4	MFV1	MFV2	MFV3	MFV4
RV1	1.00							
RV2	0.36	1.00						
RV3	0.03	0.56	1.00					
RV4	-0.11	-0.02	0.33	1.00				
MFV1	0.29	0.69	0.73	0.35	1.00			
MFV2	0.33	0.71	0.76	0.39	0.97	1.00		
MFV3	0.33	0.71	0.75	0.39	0.96	1.00	1.00	
MFV4	0.32	0.66	0.72	0.40	0.89	0.96	0.97	1.00

Japanese Yen

	RV1	RV2	RV3	RV4	MFV1	MFV2	MFV3	MFV4
RV1	1.00							
RV2	0.18	1.00						
RV3	-0.20	0.23	1.00					
RV4	-0.21	-0.09	0.36	1.00				
MFV1	0.16	0.51	0.47	0.19	1.00			
MFV2	0.11	0.54	0.58	0.25	0.90	1.00		
MFV3	0.10	0.54	0.59	0.26	0.87	1.00	1.00	
MFV4	0.05	0.47	0.56	0.25	0.69	0.92	0.95	1.00

British Pound

	RV1	RV2	RV3	RV4	MFV1	MFV2	MFV3	MFV4
RV1	1.00							
RV2	0.35	1.00						
RV3	-0.03	0.39	1.00					
RV4	-0.18	-0.20	0.22	1.00				
MFV1	0.42	0.74	0.63	0.10	1.00			
MFV2	0.43	0.74	0.62	0.10	0.98	1.00		
MFV3	0.43	0.73	0.62	0.11	0.97	1.00	1.00	
MFV4	0.40	0.71	0.60	0.11	0.93	0.98	0.99	1.00

Notes: This table reports correlations of monthly measures of realized variances (RV1, RV2, RV3, and RV4) and model-free variances (MFV1, MFV2, MFV3, and MFV4) for the three assets (S&P 500, Japanese Yen, and British Pound).

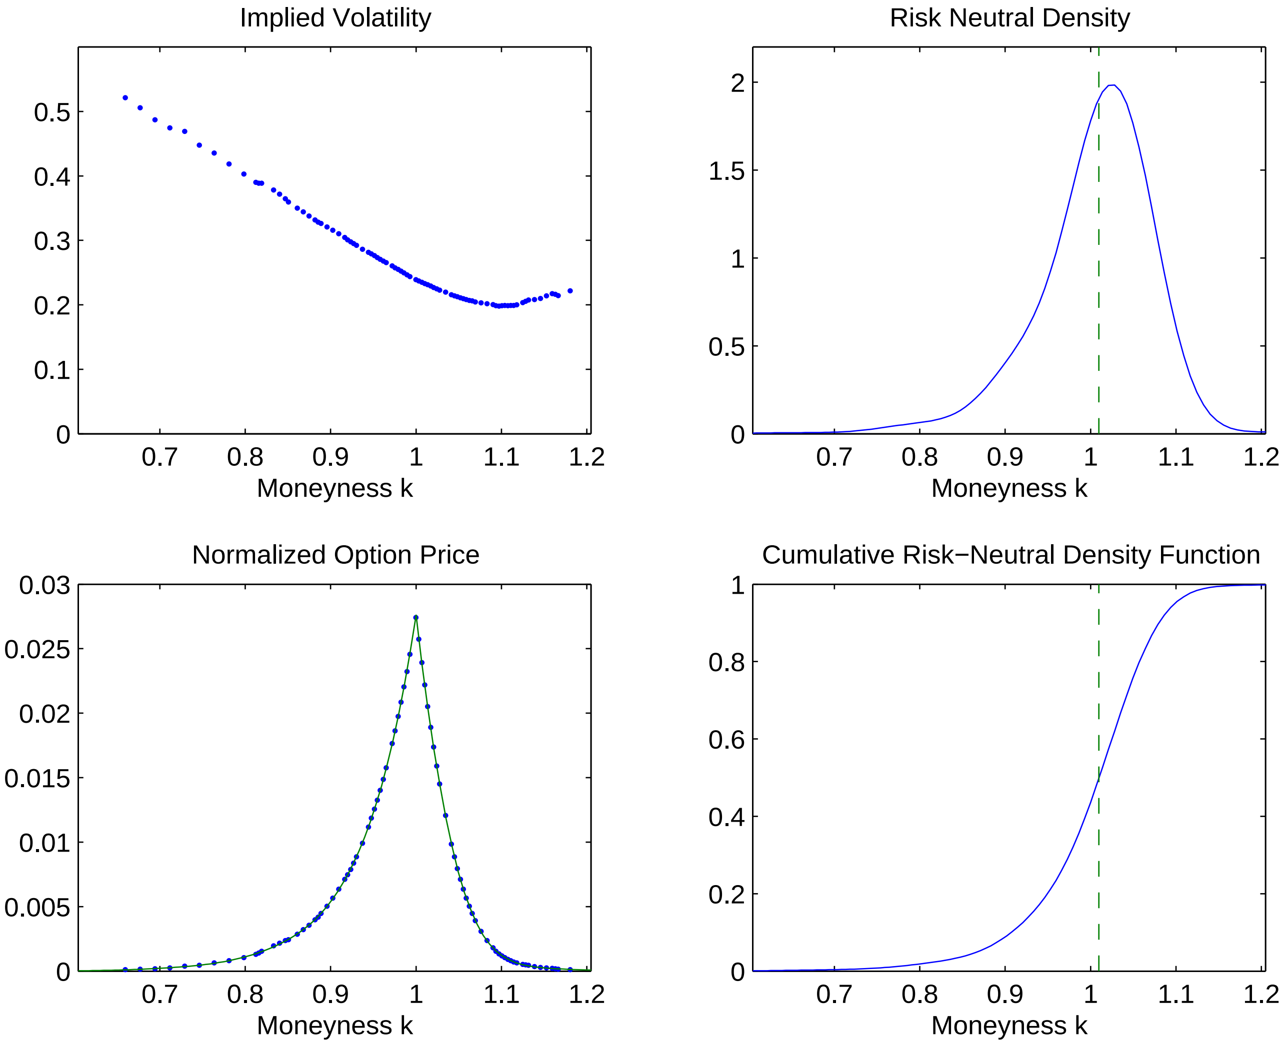

Figure 1: This figure illustrates S&P 500 option data for 04/19/2000 when $\tau = 21$ trading days. The left panels show the Black-Scholes implied volatility $IV_t(k)$ and the normalized OTM option price $M_t(k)/F_t = \min(P_t(k), C_t(k))/F_t$ for different values of moneyness $k = K/F_t$. The solid 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.

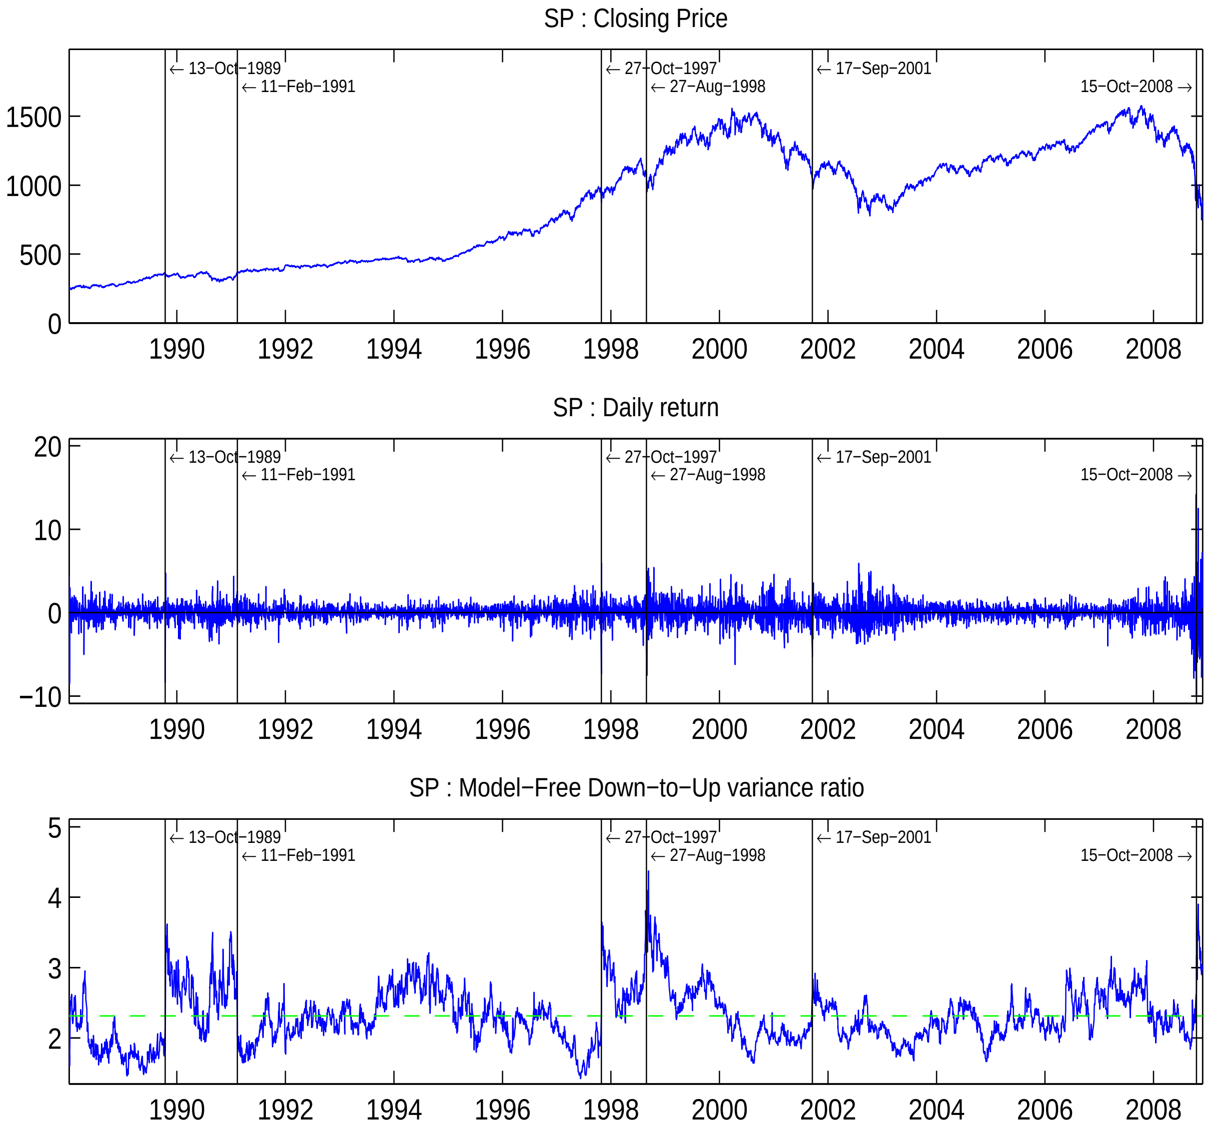

Figure 2: For S&P 500, this figure plots the daily closing futures price (the first panel), its daily return (the second panel), and the Model-Free Down-to-Up variance ratio (the third panel). The dashed line shows the sample mean of the ratio.

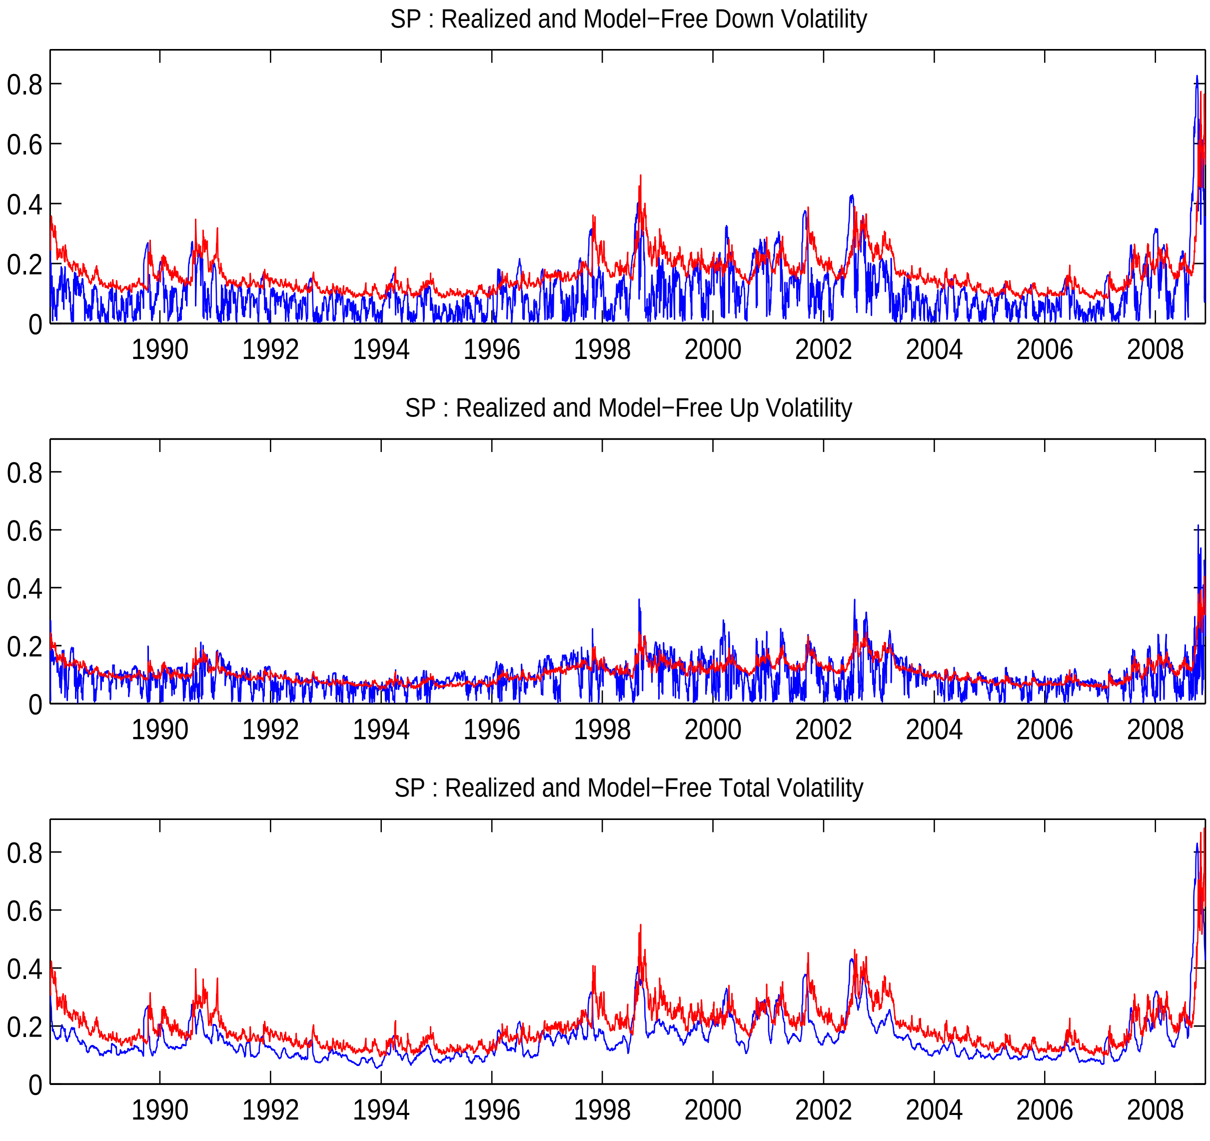

Figure 3: For S&P 500, this figure plots the 1-month realized and model-free volatilities: Down (the first panel), Up (the second panel), and Total (the third panel).

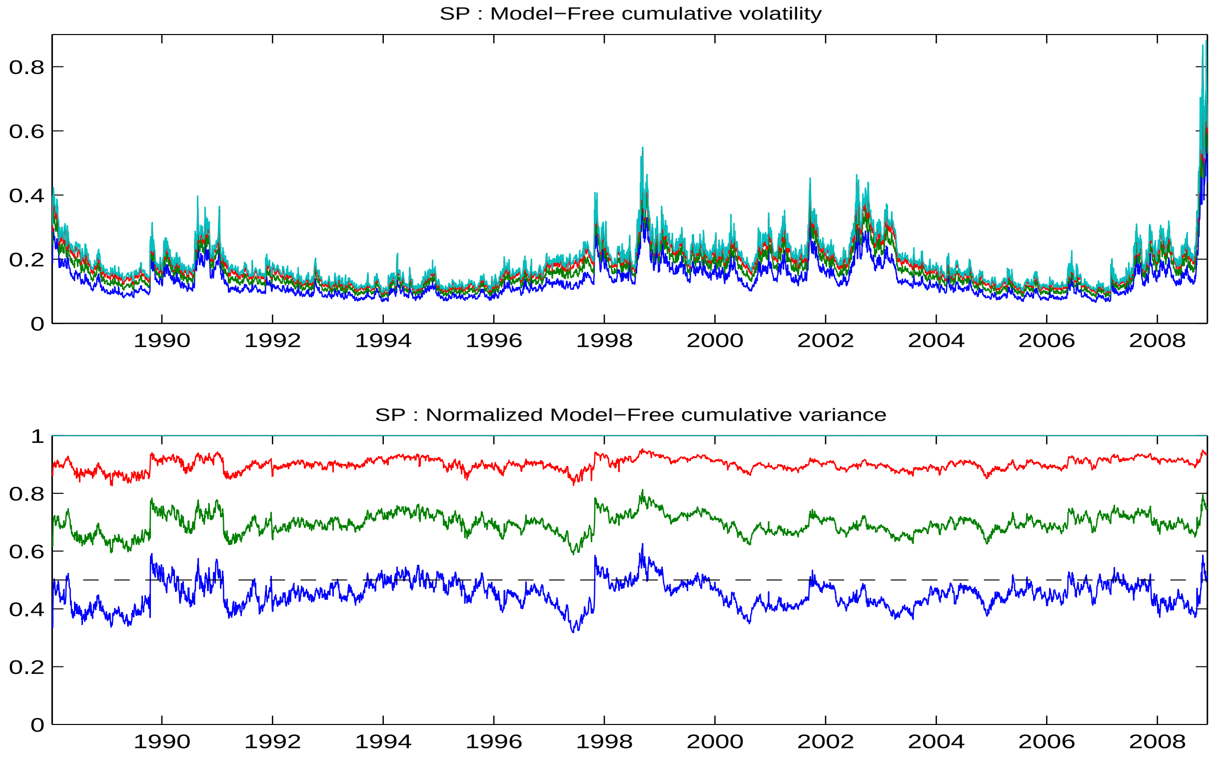

Figure 4: For S&P 500, this figure shows 1-month model-free cumulative volatilities V1-V4 (the first panel) and model-free cumulative variances scaled by the total variance MFVT.)

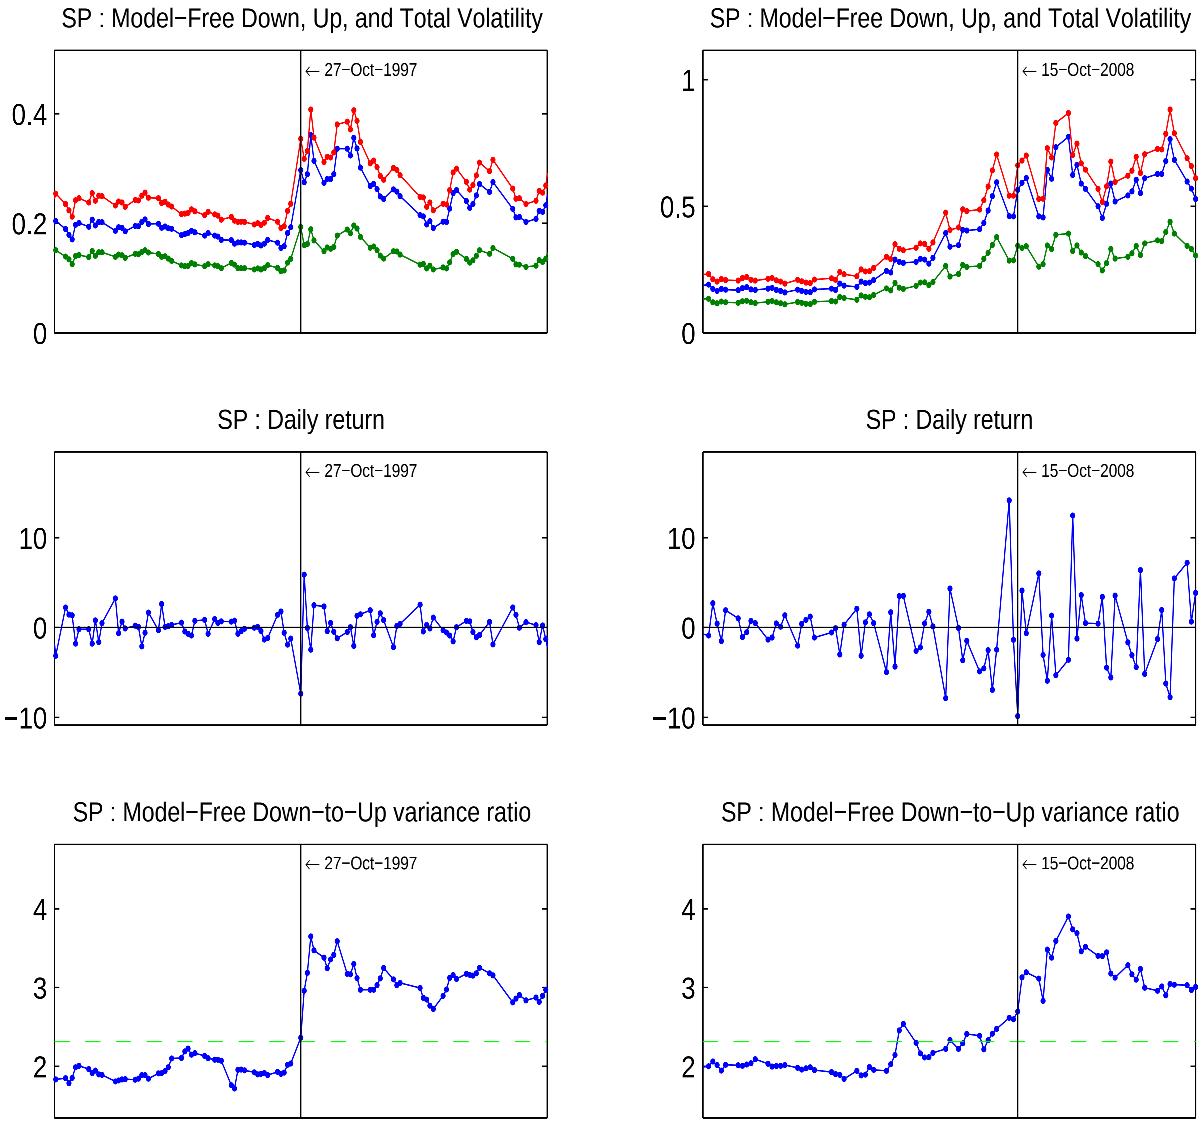

Figure 5: For S&P 500, this figure plots the Model-Free Down, Up, and Total volatility (the top panels), daily futures return (the middle panels), and the Model-Free Down-to-Up variance ratio (the bottom panels). The dashed line shows the sample mean of the ratio.

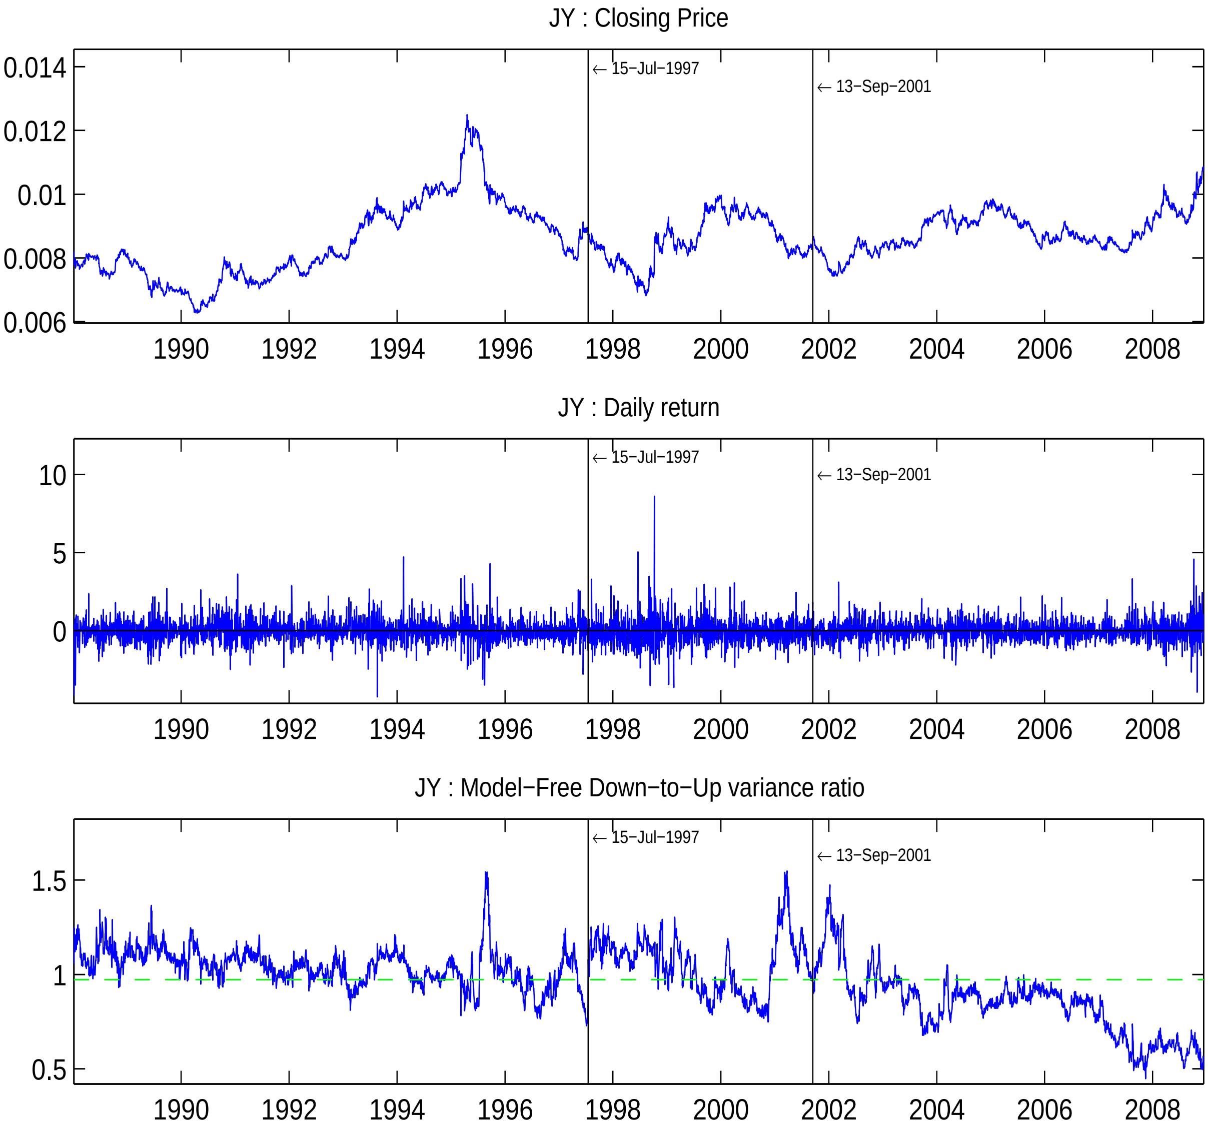

JY : Closing Price

JY : Daily return

JY : Model−Free Down−to−Up variance ratio

Figure 6: For Japanese Yen, this figure plots the closing futures price (the first panel), its daily return (the second panel), and the Model-Free Down-to-Up variance ratio (the third panel). The dashed line shows the sample mean of the ratio.

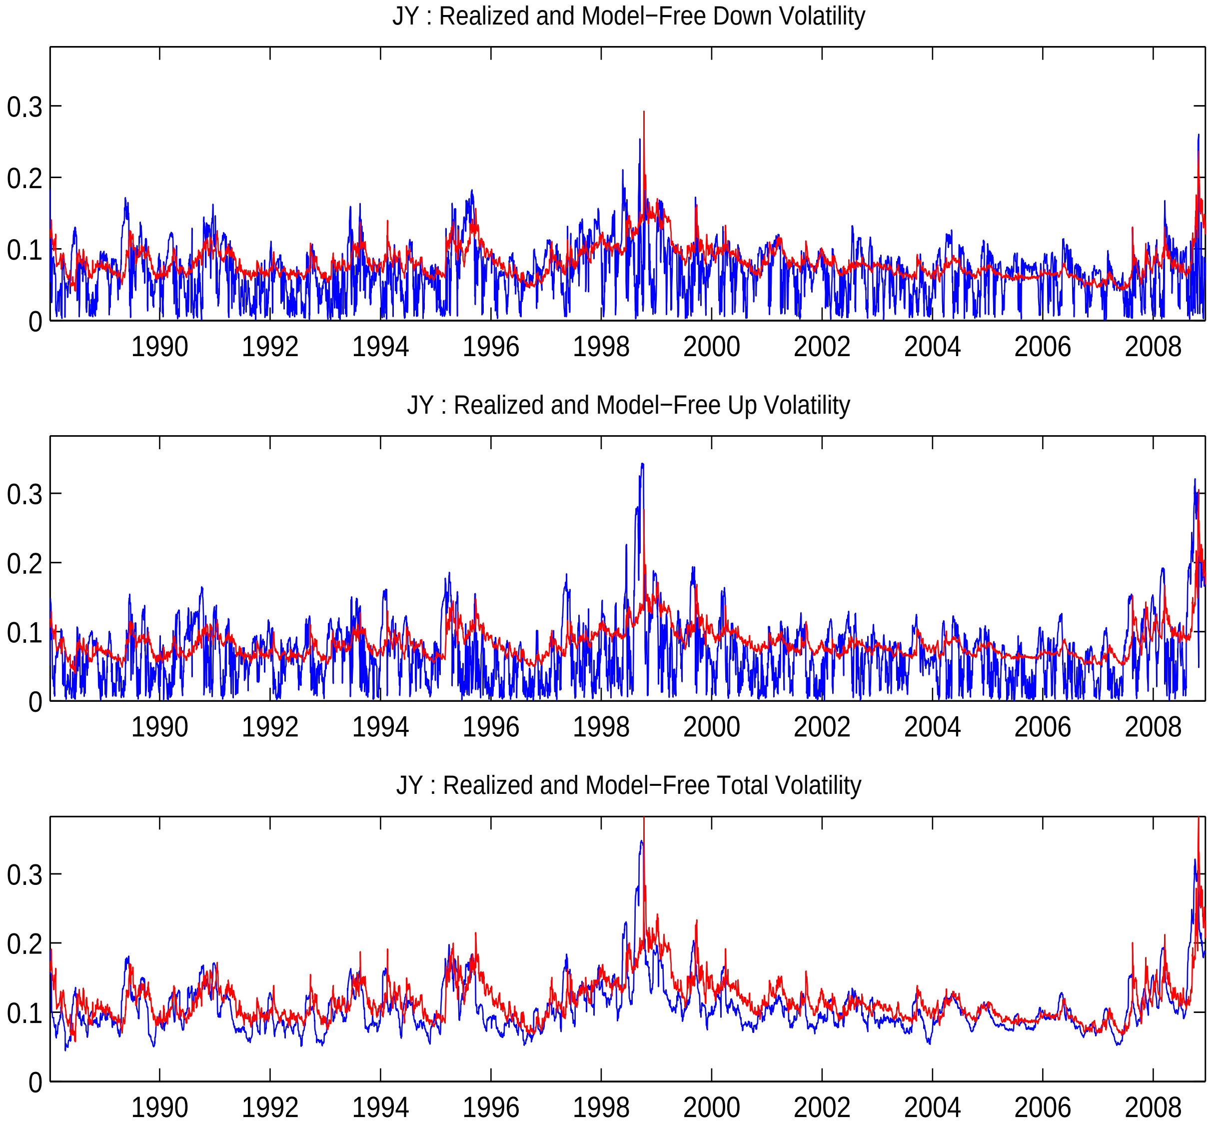
JY : Realized and Model−Free Down Volatility

JY : Realized and Model−Free Up Volatility

JY : Realized and Model−Free Total Volatility

Figure 7: For Japanese Yen, this figure plots the 1-month realized and model-free volatilities: Down (the first panel), Up (the second panel), and Total (the third panel).

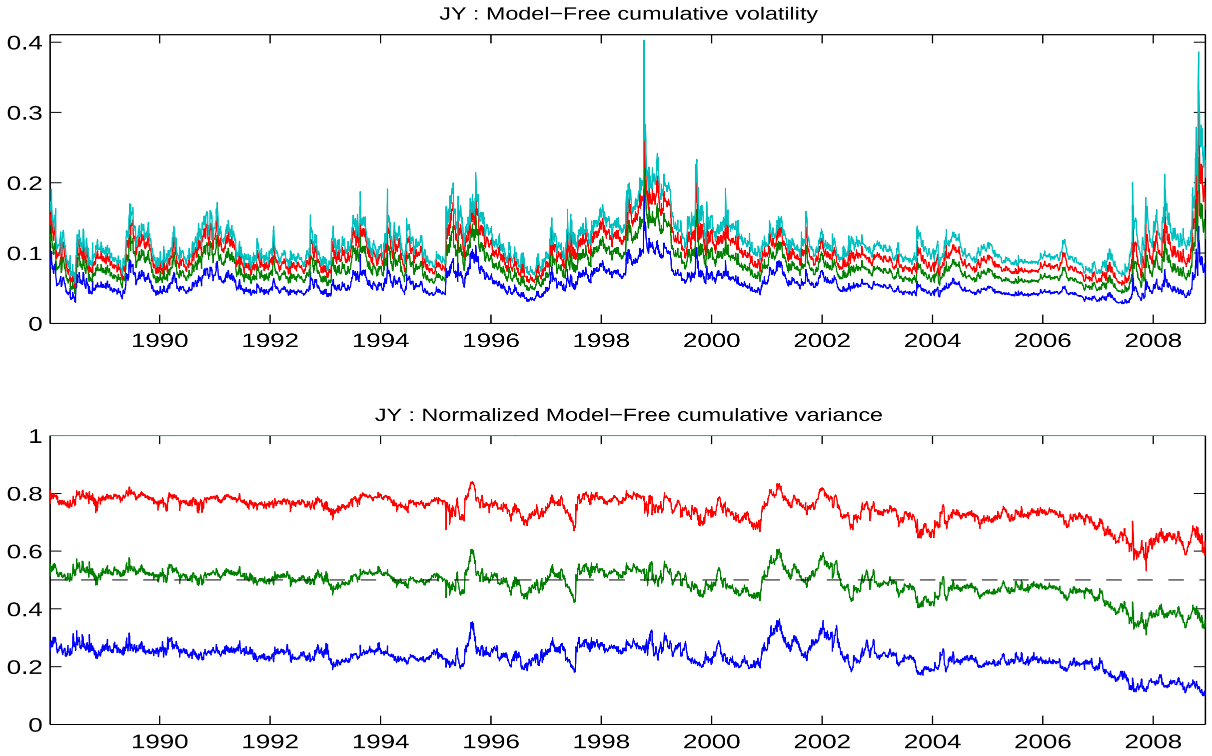

JY : Model−Free cumulative volatility

JY : Normalized Model−Free cumulative variance

Figure 8: For Japanese Yen, this figure shows 1-month model-free cumulative volatilities V1-V4 (the first panel) and model-free cumulative variances scaled by the total variance MFVT.)

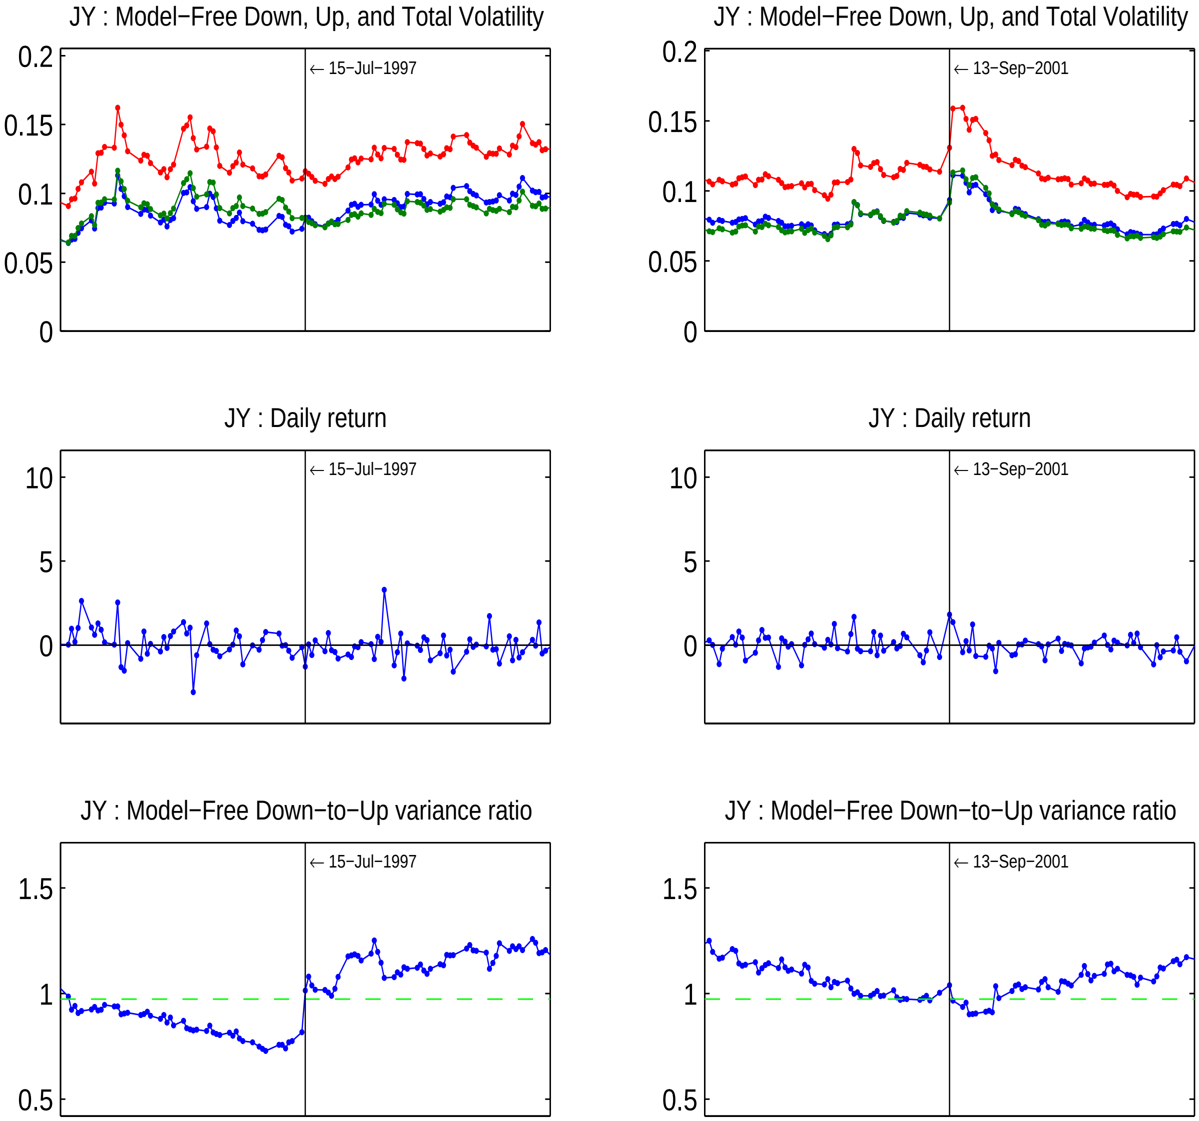
JY : Model−Free Down, Up, and Total Volatility

JY : Model−Free Down, Up, and Total Volatility

JY : Daily return

JY : Daily return

JY : Model−Free Down−to−Up variance ratio

JY : Model−Free Down−to−Up variance ratio

Figure 9: For Japanese Yen, this figure plots the Model-Free Down, Up, and Total volatility (the top panels), daily futures return (the middle panels), and the Model-Free Down-to-Up variance ratio (the bottom panels). The dashed line shows the sample mean of the ratio.

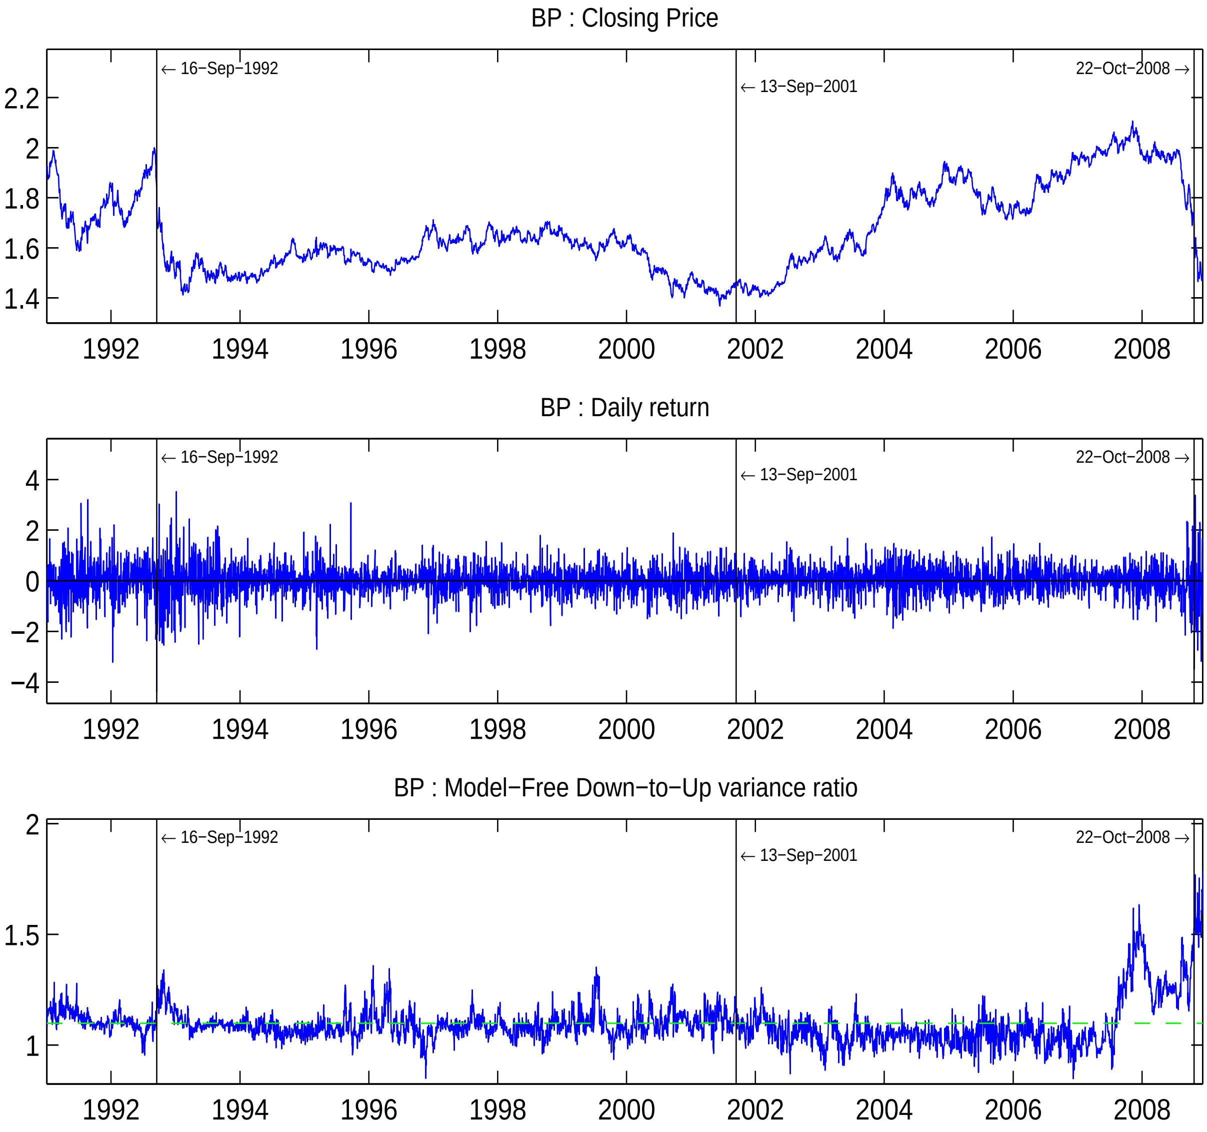

Figure 10: For British Pound, this figure plots the closing futures price (the first panel), its daily return (the second panel), and the Model-Free Down-to-Up variance ratio (the third panel). The dashed line shows the sample mean of the ratio.

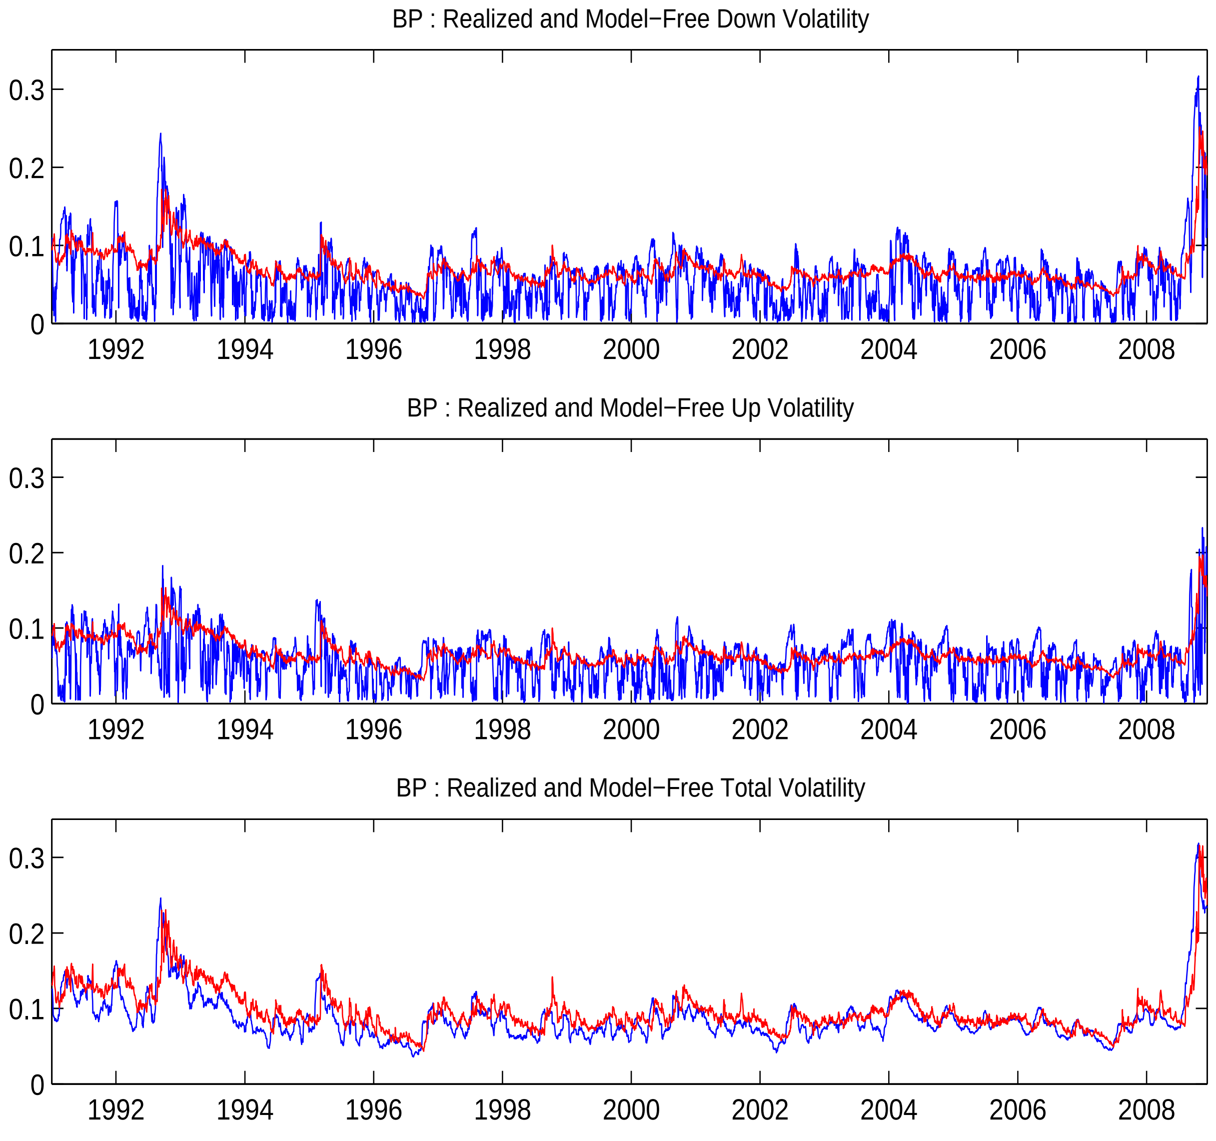

Figure 11: British Pound, this figure plots the 1-month realized and model-free volatilities: Down (the first panel), Up (the second panel), and Total (the third panel).

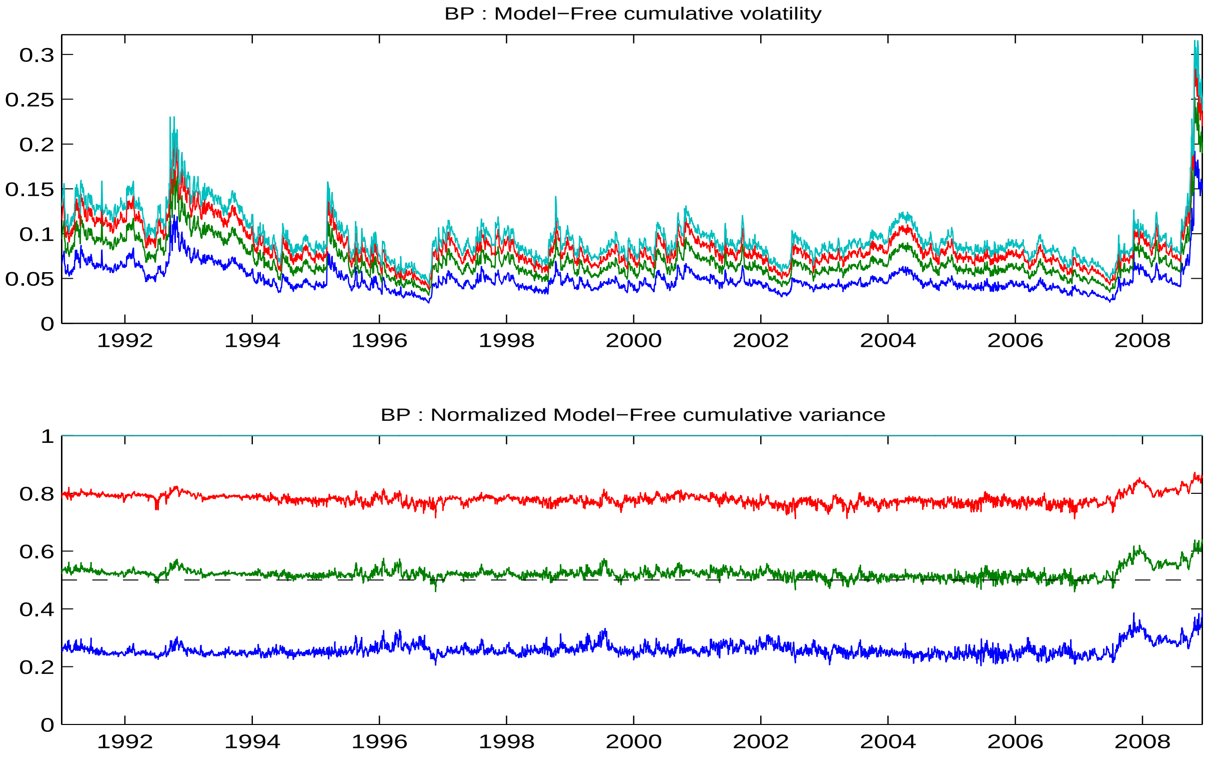
BP : Model−Free cumulative volatility

BP : Normalized Model−Free cumulative variance

Figure 12: For British Pound, this figure shows 1-month model-free cumulative volatilities V1-V4 (the first panel) and model-free cumulative variances scaled by the total variance MFVT.)

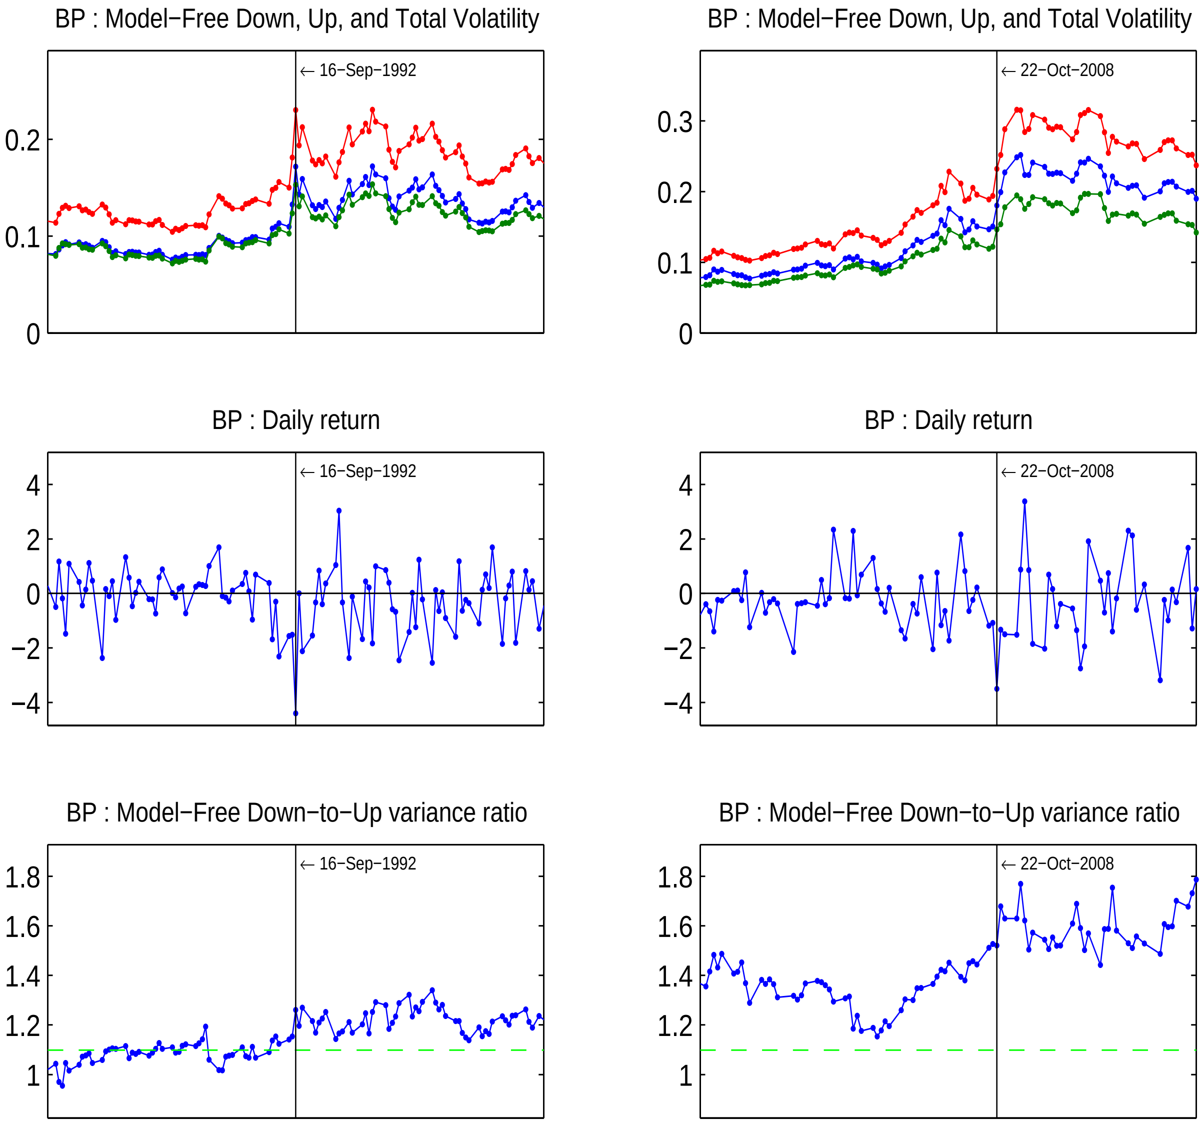

Figure 13: For British Pound, this figure plots the Model-Free Down, Up, and Total volatility (the top panels), daily futures return (the middle panels), and the Model-Free Down-to-Up variance ratio (the bottom panels). The dashed line shows the sample mean of the ratio.