Queen's Economics Department Working Paper No. 1181

The Role of Implied Volatility in Forecasting Future Realized Volatility and Jumps in Foreign Exchange, Stock, and Bond Markets

Thomas Busch
Danske Bank and CREATES

Bent Jesper Christensen
University of Aarhus and CREATES

Morten Ørregaard Nielsen
Queen's University and CREATES

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

10-2008

The Role of Implied Volatility in Forecasting Future Realized Volatility and Jumps in Foreign Exchange, Stock, and Bond Markets

Thomas Busch
Danske Bank and CREATES

Bent Jesper Christensen
University of Aarhus and CREATES

Morten Ørregaard Nielsen*
Queen's University and CREATES

June 17, 2009

Abstract

We study the forecasting of future realized volatility in the foreign exchange, stock, and bond markets from variables in the information set, including implied volatility backed out from option prices. Realized volatility is separated into its continuous and jump components, and the heterogeneous autoregressive (HAR) model is applied with implied volatility as an additional forecasting variable. A vector HAR (VecHAR) model for the resulting simultaneous system is introduced, controlling for possible endogeneity issues. We find that implied volatility contains incremental information about future volatility in all three markets, relative to past continuous and jump components, and it is an unbiased forecast in the foreign exchange and stock markets. Out-of-sample forecasting experiments confirm that implied volatility is important in forecasting future realized volatility components in all three markets. Perhaps surprisingly, the jump component is, to some extent, predictable, and options appear calibrated to incorporate information about future jumps in all three markets.

Keywords: Bipower variation, HAR, Heterogeneous Autoregressive Model, implied volatility, jumps, options, realized volatility, VecHAR, volatility forecasting.

JEL classification: C22, C32, F31, G1.

1 Introduction

In both the theoretical and empirical finance literatures, volatility is generally recognized as one of the most important determinants of risky asset values, such as exchange rates, stock and bond prices, and hence interest rates. Since any valuation procedure involves assessing the level and riskiness of future payoffs, it is particularly the forecasting of future volatility from variables in the current information set that is important for asset pricing, derivative pricing, hedging, and risk management.

A number of different variables are potentially relevant for volatility forecasting. In

*Corresponding author. Address: Department of Economics, Dunning Hall Room 307, 94 University Avenue, Queen's University, Kingston, Ontario K7L 3N6, Canada. Email: mon@econ.queensu.ca

the present paper, we include derivative prices and investigate whether implied volatilities (IV) backed out from options on foreign currency futures, stock index futures, or Treasury bond (T-bond) futures contain incremental information when assessed against volatility forecasts based on high-frequency (5-minute) current and past returns on exchange rates, stock index futures, and T-bond futures, respectively.

Andersen, Bollerslev, Diebold & Labys (2003) and Andersen, Bollerslev & Meddahi (2004) show that simple reduced form time series models for realized volatility (RV) outperform commonly used GARCH and related stochastic volatility models in forecasting future volatility. In recent work, Barndorff-Nielsen & Shephard (2004, 2006) derive a fully nonparametric separation of the continuous sample path (C) and jump (J) components of RV. Applying this technique, Andersen, Bollerslev & Diebold (2007) extend results of Andersen et al. (2003) and Andersen et al. (2004) by using past C and J as separate regressors when forecasting volatility. They show that the two components play very different roles in forecasting, and that significant gains in performance are achieved by separating them. While C is strongly serially correlated, J is distinctly less persistent, and almost not forecastable, thus clearly indicating separate roles for C and J in volatility forecasting.

In this paper, we study high-frequency (5-minute) returns to the \$/DM exchange rate, S&P 500 futures, and 30 year T-bond futures, as well as monthly prices of associated futures options. Alternative volatility measures are computed from the two separate data segments, i.e., RV and its components from high-frequency returns and IV from option prices. IV is widely perceived as a natural forecast of integrated volatility over the remaining life of the option contract under risk-neutral pricing. It is also a relevant forecast in a stochastic volatility setting even if volatility risk is priced, and it should get a coefficient below (above) unity in forecasting regressions in case of a negative (positive) volatility risk premium (Bollerslev & Zhou (2006)). Since options expire at a monthly frequency, we consider the forecasting of one-month volatility measures. The issue is whether IV retains incremental information about future integrated volatility when assessed against realized measures (RV, C, J) from the previous month. The methodological contributions of the present paper are to use high-frequency data and recent statistical techniques for the realized measures, and to allow these to have different impacts at different frequencies, when constructing the return-based forecasts that IV is assessed against. These innovations ensure that IV is put to a harder test than in previous literature when comparing forecasting performance.

The idea of allowing different impacts at different frequencies arises since realized measures covering the entire previous month very likely are not the only relevant yardsticks. Squared returns nearly one month past may not be as informative about future volatility as squared returns that are only one or a few days old. To address this issue, we apply the heterogeneous autoregressive (HAR) model proposed by Corsi (2009) for RV analysis and extended by Andersen et al. (2007) to include the separate C and J components of total realized volatility ($RV = C+J$) as regressors. In the HAR framework, we include IV from option prices as an additional regressor, and also consider separate forecasting of both C and J individually. As an additional contribution, we introduce a vector heterogeneous autoregressive (labeled VecHAR) model for joint modeling of IV, C, and J. Since IV is the new variable added in our study, compared to the RV literature, and since it may potentially be measured with error stemming from non-synchronicity between sampled option prices and corresponding futures prices, bid-ask spreads, model error, etc., we take special care in handling this variable. The simultaneous VecHAR analysis controls for possible endogeneity issues in the forecasting equations, and allows testing interesting cross-equation restrictions.

Based on in-sample Mincer & Zarnowitz (1969) forecasting regressions, we show that IV contains incremental information relative to both C and J when forecasting subsequent RV in all three markets. Furthermore, in the foreign exchange and stock markets, IV is an unbiased forecast. Indeed, it completely subsumes the information content of the daily, weekly, and monthly high-frequency realized measures in the foreign exchange market. Moreover, out-of-sample forecasting evidence suggests that IV should be used alone when forecasting monthly RV in all three markets. The mean absolute out-of-sample forecast error increases if any RV components are included in constructing the forecast.

Using the HAR methodology for separate forecasting of C and J, our results show that IV has predictive power for each. Forecasting monthly C is very much like forecasting RV itself. The coefficient on IV is slightly smaller, but in-sample qualitative results on which variables to include are identical. The out-of-sample forecasting evidence suggests that IV again should be used alone in the foreign exchange and stock markets, but that it should be combined with realized measures in the bond market. Perhaps surprisingly, even the jump component is, to some extent, predictable, and IV contains incremental information about future jumps in all three markets.

The results from the VecHAR model reinforce the conclusions. In particular, when forecasting C in the foreign exchange market, IV completely subsumes the information content of all realized measures. Out-of-sample forecasting performance is about unchanged for J but improves for C in all markets by using the VecHAR model, relative to comparable univariate specifications. The VecHAR system approach allows testing cross-equation restrictions, the results of which support the finding that IV is a forecast of total realized volatility $RV = C + J$, indeed an unbiased forecast in the foreign exchange and stock markets.

In the previous literature, a number of authors have included IV in forecasting regressions, and most have found that it contains at least some incremental information, although there is mixed evidence on its unbiasedness and efficiency.[1] None of these studies has investigated whether the finding of incremental information in IV holds up when separating C and J computed from high-frequency returns, or when including both daily, weekly, and monthly realized measures in HAR-type specifications. An interesting alternative to using individual option prices might have been to use model-free implied volatilities as in Jiang & Tian (2005). However, Andersen & Bondarenko (2007) find that these are dominated by the simpler Black-Scholes implied volatilities in terms of forecasting power.

[1] See, e.g., Jorion (1995), Xu & Taylor (1995), Covrig & Low (2003), and Pong, Shackleton, Taylor & Xu (2004) on the foreign exchange market, Day & Lewis (1992), Canina & Figlewski (1993), Lamoureux & Lastrapes (1993), Christensen & Prabhala (1998), Fleming (1998), and Blair, Poon & Taylor (2001) on the stock market, and Amin & Morton (1994) on the bond market.

The remainder of the paper is laid out as follows. In the next section we briefly describe realized volatility and the nonparametric identification of its separate continuous sample path and jump components. In section 3 we discuss the derivative pricing model. Section 4 describes our data. In section 5 the empirical results are presented, and section 6 concludes.

2 The Econometrics of Jumps

We assume that the logarithm of the asset price, $p(t)$, follows the general stochastic volatility jump diffusion model

$$dp(t) = \mu(t)\,dt + \sigma(t)\,dw(t) + \kappa(t)\,dq(t), \quad t \geq 0. \tag{1}$$

The mean $\mu(\cdot)$ is assumed continuous and locally bounded, the instantaneous volatility $\sigma(\cdot) > 0$ is càdlàg, and $w(\cdot)$ is the driving standard Brownian motion. The counting process $q(t)$ is normalized such that $dq(t) = 1$ corresponds to a jump at time t and $dq(t) = 0$ otherwise. Hence, $\kappa(t)$ is the random jump size at time t if $dq(t) = 1$. The intensity of the arrival process for jumps, $\lambda(t)$, is possibly time-varying, but does not allow infinite activity jump processes. Note that the leverage effect is accommodated in (1) through possible dependence between $\sigma(\cdot)$ and $w(\cdot)$, see Barndorff-Nielsen, Graversen, Jacod & Shephard (2006) and Barndorff-Nielsen, Shephard & Winkel (2006).

The quadratic variation $[p](t)$ is defined for any semimartingale by

$$[p](t) = p\lim \sum_{j=1}^{K} \left(p(s_j) - p(s_{j-1})\right)^2, \tag{2}$$

where $0 = s_0 < s_1 < ... < s_K = t$ and the limit is taken for $\max_j |s_j - s_{j-1}| \to 0$ as $K \to \infty$. In the model (1), we have in wide generality

$$[p](t) = \int_0^t \sigma^2(s)\,ds + \sum_{j=1}^{q(t)} \kappa^2(t_j), \tag{3}$$

where $0 \leq t_1 < t_2 < ...$ are the jump times. In (3), quadratic variation is decomposed as integrated volatility plus the sum of squared jumps through time t.

Assume that $M+1$ evenly spaced intra-period observations for period t are available on the log-price $p_{t,j}$. The continuously compounded intra-period returns are

$$r_{t,j} = p_{t,j} - p_{t,j-1}, \quad j = 1, ..., M, \quad t = 1, ..., T, \tag{4}$$

where T is the number of periods in the sample. Realized volatility for period t is given by the sum of squared intra-period returns,

$$RV_t = \sum_{j=1}^{M} r_{t,j}^2, \quad t = 1, ..., T. \tag{5}$$

Following Barndorff-Nielsen & Shephard (2004, 2006), the nonparametric separation of the continuous sample path and jump components of quadratic variation in (3) can be done through the related bipower and tripower variation measures. The staggered (skip-k, with $k \geq 0$) realized bipower variation is defined as

$$BV_t = \mu_1^{-2} \frac{M}{M-(k+1)} \sum_{j=k+2}^{M} |r_{t,j}| \, |r_{t,j-k-1}| \,, \quad t = 1, ..., T, \tag{6}$$

where $\mu_1 = \sqrt{2/\pi}$. In theory, a higher value of M improves precision of the estimators, but in practice it also makes them more susceptible to market microstructure effects, such as bid-ask bounces, stale prices, measurement errors, etc., introducing artificial (typically negative) serial correlation in returns, see, e.g., Hansen & Lunde (2006) and Barndorff-Nielsen & Shephard (2007). Huang & Tauchen (2005) show that staggering (i.e., setting $k \geq 1$) mitigates the resulting bias in (6), since it avoids the multiplication of the adjacent returns $r_{t,j}$ and $r_{t,j-1}$ that by (4) share the log-price $p_{t,j-1}$ in the non-staggered (i.e., $k = 0$) version of (6). Further, staggered realized tripower quarticity is

$$TQ_t = \mu_{4/3}^{-3} \frac{M^2}{M-2(k+1)} \sum_{j=2k+3}^{M} |r_{t,j}|^{4/3} \, |r_{t,j-k-1}|^{4/3} \, |r_{t,j-2k-2}|^{4/3} \,, \quad t = 1, ..., T, \tag{7}$$

where $\mu_{4/3} = 2^{2/3}\Gamma(7/6)/\Gamma(1/2)$. We follow Huang & Tauchen (2005) and use $k = 1$ in (6) and (7) in our empirical work. The choice of k has no impact on asymptotic results.

Combining (2) and (5), RV_t is by definition a consistent estimator of the per-period increment $[p](t) - [p](t-1)$ to quadratic variation as $M \to \infty$. At the same time, BV_t is consistent for the integrated volatility portion of the increment,

$$BV_t \to_p \int_{t-1}^{t} \sigma^2(s)\, ds \quad \text{as } M \to \infty, \tag{8}$$

as shown by Barndorff-Nielsen & Shephard (2004) and Barndorff-Nielsen, Shephard & Winkel (2006). It follows that the difference $RV_t - BV_t$ converges to the sum of squared jumps that have occurred during the period. Of course, it may be non-zero in finite samples due to sampling variation even if no jump occurred during period t, so a notion of a "significant jump component" is needed. Following e.g. Huang & Tauchen (2005) and Andersen et al. (2007), we apply the (ratio) test statistic

$$Z_t = \sqrt{M} \frac{(RV_t - BV_t)RV_t^{-1}}{\left((\mu_1^{-4} + 2\mu_1^{-2} - 5) \max\{1, TQ_t BV_t^{-2}\} \right)^{1/2}}. \tag{9}$$

In the absence of jumps, $Z_t \to_d N(0,1)$ as $M \to \infty$, and large positive values indicate that jumps occurred during period t. Huang & Tauchen (2005) show that market microstructure noise may bias the test against finding jumps, but also that staggering alleviates the bias.

The (significant) jump component of realized volatility is now

$$J_t = I_{\{Z_t > \Phi_{1-\alpha}\}} (RV_t - BV_t), \quad t = 1, ..., T, \tag{10}$$

where $I_{\{A\}}$ is the indicator for the event A, $\Phi_{1-\alpha}$ the $100(1-\alpha)\%$ point in the standard normal distribution, and α the significance level. When $I_{\{Z_t > \Phi_{1-\alpha}\}} = 1$, J_t is excess realized volatility above bipower variation, and hence attributable to jumps in prices. The continuous component of quadratic variation is estimated by the remainder of RV_t,

$$C_t = RV_t - J_t, \quad t = 1, ..., T. \tag{11}$$

This way, C_t equals RV_t if there is no significant jump during period t, and BV_t if there is, i.e., $C_t = I_{\{Z_t \leq \Phi_{1-\alpha}\}} RV_t + I_{\{Z_t > \Phi_{1-\alpha}\}} BV_t$. For any standard significance level $\alpha < 1/2$, both J_t and C_t from (10) and (11) are non-negative because $\Phi_{1-\alpha}$ is. Consistency of each component as estimators of the corresponding components of quadratic variation, i.e.,

$$C_t \to_p \int_{t-1}^{t} \sigma^2(s)\, ds \text{ and } J_t \to_p \sum_{j=q(t-1)+1}^{q(t)} \kappa^2(t_j),$$

may be achieved by letting $\alpha \to 0$ and $M \to \infty$ simultaneously. Hence, this high-frequency data approach allows for period-by-period nonparametric consistent estimation of both components of quadratic variation in (3).

3 Derivative Pricing Model

For the construction of implied volatility, we let c denote the call option price, X the exercise or strike price, τ the time to expiration of the option, F the price of the underlying futures contract with delivery date Δ periods after option expiration, and r the riskless interest rate. We use the futures option pricing formula, see Bates (1996a, 1996b),

$$c(F, X, \tau, \Delta, r, \sigma^2) = e^{-r(\tau+\Delta)}[F\Phi(d) - X\Phi(d - \sqrt{\sigma^2\tau})], \quad d = \frac{\ln(F/X) + \frac{1}{2}\sigma^2\tau}{\sqrt{\sigma^2\tau}}, \tag{12}$$

where $\Phi(\cdot)$ is the standard normal c.d.f. and σ the futures return volatility. The case $\Delta = 0$ (no delivery lag) corresponds to the well-known Black (1976) and Garman & Kohlhagen (1983) futures option formula. For general $\Delta > 0$, regarding the futures contract as an asset paying a continuous dividend yield equal to the riskless rate r, the asset price in the standard Black & Scholes (1973) and Merton (1973) formula is replaced by the discounted futures price $e^{-r(\tau+\Delta)}F$. Jorion (1995) applied (12) with $\Delta = 0$ to the currency option market, whereas Bates (1996a, 1996b) used delivery lags Δ specific to the Philadelphia Exchange (PHLX) and the Chicago Mercantile Exchange (CME), respectively.

We consider serial \$/DM and S&P 500 futures options with monthly expiration cycle traded at the CME, and equivalent T-bond futures options traded at the Chicago Board of Trade (CBOT). The contract specifications do not uniquely identify the particular T-bond serving as underlying asset for the bond futures, requiring merely that it does not mature and is not callable for at least 15 years from the first day of the delivery month of the underlying futures. The delivery month of the underlying futures contracts follows a quarterly (March) cycle, with delivery date on the third Wednesday of the month for currency and bond futures, and the third Friday for stock index futures. The options expire two Fridays prior to the third Wednesday of each month in the currency case, on the last Friday followed by at least two business days in the month in the bond case, and on the third Friday in the stock case, except every third month where it is shifted to the preceding Thursday to avoid "triple witching hour" problems associated with simultaneous maturing of the futures, futures options, and index options. Upon exercise, the holder of the option receives a position at the strike X in the futures, plus the intrinsic value $F-X$ in cash, on the following trading day, so the delivery lag is $\Delta = 3/365$ (from Friday to Monday), except $\Delta = 1/365$ (Thursday to Friday) every third month in the stock case. Finally, following French (1984), τ is measured in trading days when used with volatilities ($\sigma^2\tau$ in (12)) and in calender days when concerning interest rates (in $r(\tau+\Delta)$).

Given observations on the option price c and the variables F, X, τ, Δ, and r, an implied volatility (IV) estimate in variance form can be backed out from (12) by numerical inversion of the nonlinear equation $c = c(F, X, \tau, \Delta, r, IV)$ with respect to IV. In our empirical work, IV measured one month prior to expiration is used as a forecast of subsequent RV and its components C and J (section 2), measured from high-frequency returns over the remaining life of the option, i.e., the one-month interval ending at expiration. For stocks and bonds, these are returns on the futures, i.e., the underlying asset. In the currency case, we use returns to the $/DM spot exchange rate. Differences between this and the futures rate underlying the option stem mainly from the interest differential in the interest rate parity

$$\ln F = p + (r_{\$} - r_{DM})\tau \tag{13}$$

from international finance and should be slight. Condition (13) is exact for constant interest rates, since then forward and futures prices coincide (Cox, Ingersoll & Ross (1981)), and an approximation for stochastic rates. Indeed, under (13), the Garman & Kohlhagen (1983) spot exchange rate option pricing formula reduces to (12). This European style formula is here applied to American style options since early exercise premia are very small for short-term, at-the-money (ATM, $X = F$) futures options, as noted by Jorion (1995).

Although (12) is used as a common standard among practitioners and in the empirical literature, its derivation does not accommodate jumps, and hence J may not be forecast very well by IV backed out from this formula. On the other hand, it is consistent with a time-varying volatility process for a continuous sample path futures price process, suggesting that IV should have better forecasting power for C. In fact, our empirical results below show that IV can predict both C and J, although it is confirmed that J is the most difficult component to predict. The findings suggest that in practice, option prices are, at least to some extent, calibrated to incorporate the possibility of future jumps. For our analysis, this reduces the empirical need to invoke a more general option pricing formula explicitly allowing for jumps. Such an extension would entail estimation of additional parameters, including prices of volatility and jump risk, which would be a considerable complication. If anything, the results would only reveal that option prices contain even more information than that reflected in our IV measure.

4 Data Description

Serial futures options with monthly expiration cycle were introduced in January 1987 (month where option price is sampled—expiration is the following month) in the $/DM market and in October 1990 for 30 year T-bonds. Our option price data cover the period from inception through May 1999 for \$/DM and through November 2002 for bonds, and from January 1990 through November 2002 for S&P 500 futures options. We use open auction closing prices of one-month ATM calls obtained from the Commodity Research Bureau, recorded on the Tuesday after the expiration date of the preceding option contract. The US Eurodollar deposit one-month middle rate from Datastream is used for the risk-free rate r. The final samples are time series of length n of annualized IV measures from (12), covering nonoverlapping one-month periods, with $n = 149$ for the currency market[2], 155 for the stock market, and 146 for the bond market.

For RV and its separate components we use the same data as Andersen et al. (2007). These are based on five-minute observations on \$/DM spot exchange rates, S&P 500 futures prices, and T-bond futures prices. There is a total of 288 high-frequency returns per day ($r_{t,j}$ from (4)) for the currency market, 97 per day for the stock market, and 79 for the bond market. We use the nonparametric procedure from section 2 to construct monthly realized volatility measures (in annualized terms) covering exactly the same periods as the IV estimates, so each of the n months has M about 6,336 (288 returns per day for approximately 22 trading days) for the foreign exchange market, 2,134 (22×97) for the stock market, and 1,738 (22×79) for the bond market. As suggested by Andersen et al. (2007), a significance level of $\alpha = 0.1\%$ is used to detect jumps and construct the series for J and C from (10) and (11). We find significant jumps in 148 out of the $n = 149$ months in the foreign exchange market, in 120 of the 155 months in stock market, and in 138 of the 146 months in the bond market. Thus, jumps are non-negligible in all three markets.

If implied volatility were given by the conditional expectation of future realized volatility, we would expect that RV and IV had equal unconditional means, and RV higher unconditional standard deviation in the time series than IV. This pattern is confirmed in the foreign exchange and stock markets, where both RV and C have higher sample standard deviations (.007 and .006 in the foreign exchange market, .032 and .027 in the stock market) than IV (.005 in the foreign exchange market and .024 in the stock market), and almost in the bond market, where RV, C, and IV have about the same standard deviation (.003). The unconditional sample mean of IV is slightly higher than that of RV in the stock (.032 vs .029) and bond (.009 vs .007) markets, possibly reflecting a negative price of volatility risk (c.f. Bollerslev & Zhou (2006)), an early exercise premium, or the overnight closing period in the stock and bond markets. The opposite is the case for the foreign exchange market (.012 vs .013), where there is round-the-clock trading.

Time series plots of the four monthly volatility measures are exhibited in Figure 1, where Panel A is for the foreign exchange market, Panel B the stock market, and Panel C the bond market. In the foreign exchange and stock markets, the continuous component C of realized volatility is close to RV itself. The new variable in our analysis, implied volatility IV, is also close to RV, but not as close as C. In the bond market (Panel C), C is below RV, and IV

[2] Trading in \$/DM options declined near the introduction of the Euro, and for January 1999 no one-month currency option price is available, even though quarterly contract prices are. An IV estimate is constructed by linear interpolation between IV for December 1998 and February 1999.

Figure 1: Time series plots of monthly volatility measures

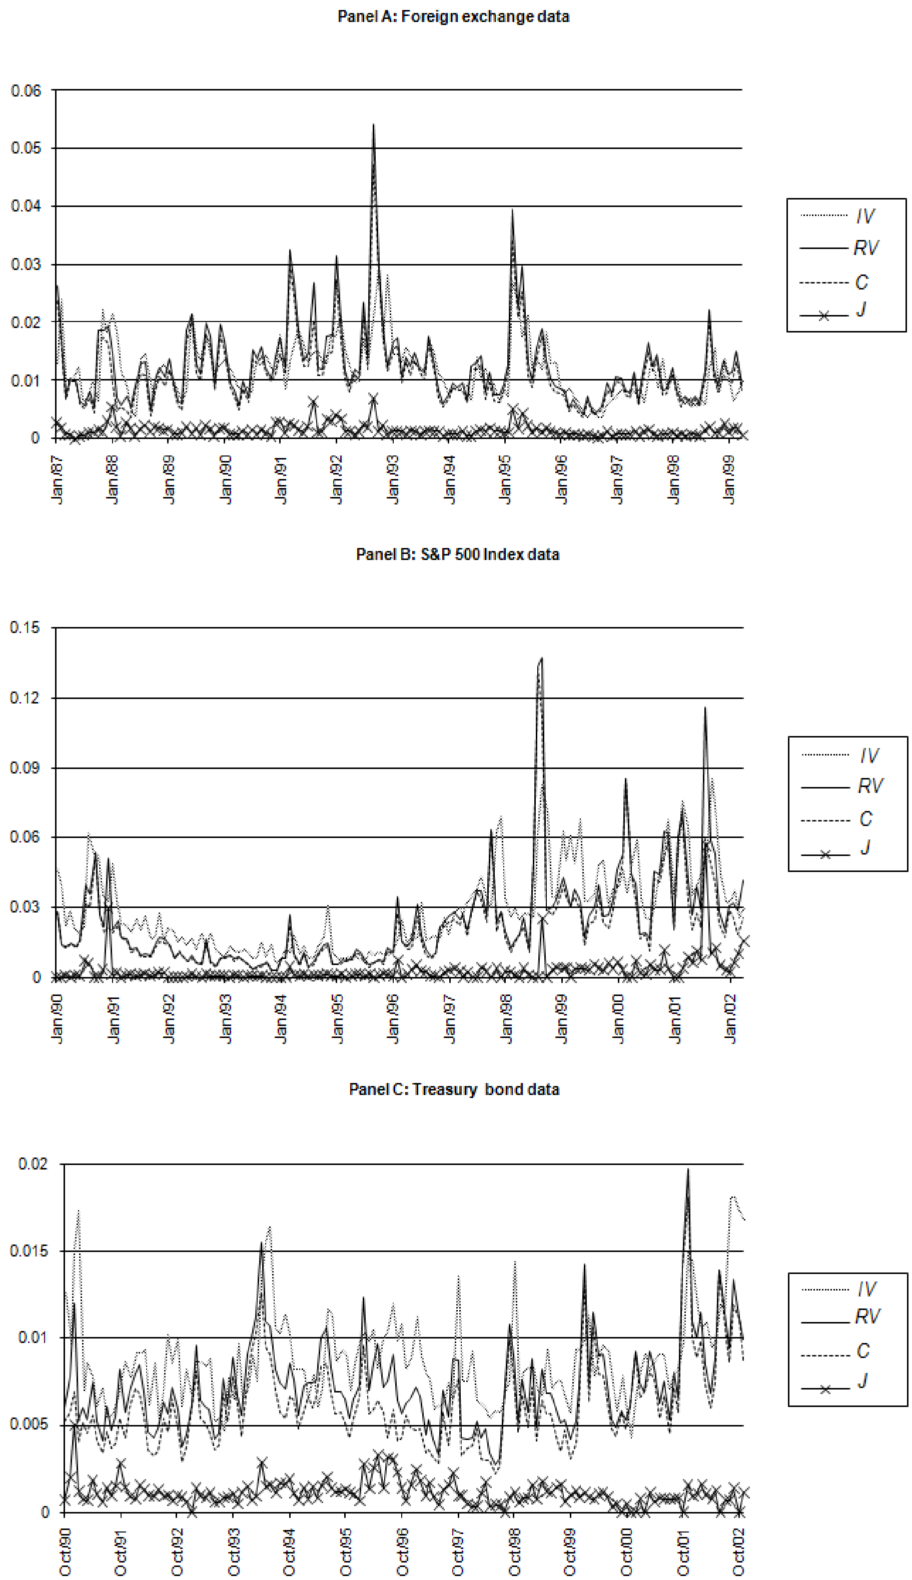

hovers above both. In all three markets, the jump component J appears to exhibit less serial dependence than the other volatility measures, consistent with Andersen et al. (2007). This is evidence of the importance of analyzing the continuous and jump components separately.

5 Econometric Models and Empirical Results

In this section we study the relation between realized volatility together with its disentangled components and implied volatility from the associated option contract. We apply the Heterogeneous Autoregressive (HAR) model in our setting with implied volatility, and we introduce a multivariate extension. Each table has results for the foreign exchange market in Panel A, the stock market in Panel B, and the T-bond market in Panel C.

5.1 Heterogeneous Autoregressive (HAR) Model

In forecasting future realized volatility, it may be relevant to place more weight on recent squared returns than on those from the more distant past. This is done in a parsimonious fashion in the HAR model of Corsi (2009). When applying it to RV itself, we denote the model HAR-RV, following Corsi (2009). When separating the RV regressors into their C and J components, we denote the model HAR-RV-CJ, following Andersen et al. (2007).

We modify the HAR-RV-CJ model in three directions. First, we include implied volatility (IV) as an additional regressor and abbreviate the resulting model HAR-RV-CJIV. Secondly, in the following subsection the HAR approach is applied to separate forecasting of C and J, rather than total realized volatility $RV = C + J$, and we denote the corresponding models HAR-C-CJIV and HAR-J-CJIV, respectively. Thirdly, Andersen et al. (2007) estimate HAR models with the regressand sampled at overlapping time intervals, e.g., monthly RV is sampled at the daily frequency, causing serial correlation in the error term. This does not necessarily invalidate the parameter estimates, although an adjustment must be made to obtain correct standard errors. However, options expire according to a monthly cycle (section 3), and the analysis in Christensen & Prabhala (1998) suggests that use of overlapping data may lead to erroneous inferences in a setting involving both IV and RV. Thus, in all our regression specifications, see e.g. (14) below, the regressand and regressors cover nonoverlapping time intervals.

The h-day realized volatility in annualized terms is

$$RV_{t,t+h} = 252h^{-1}\left(RV_{t+1} + RV_{t+2} + \ldots + RV_{t+h}\right).$$

Henceforth, we use t to indicate trading days. Thus, RV_t is now the daily realized volatility for day t from (5), while $RV_{t,t+h}$ is a daily ($h = 1$), weekly ($h = 5$), or monthly ($h = 22$) realized volatility, and similarly for the continuous ($C_{t,t+h}$) and jump ($J_{t,t+h}$) components, where the jump test Z_t from (9) is implemented using h-day period lengths for RV, BV, TQ, and M. Note that $RV_{t,t+1} = 252RV_{t+1}$, and under stationarity $E\left[RV_{t,t+h}\right] = 252E\left[RV_{t+1}\right]$ for all h. Each monthly realized measure is constructed using a value of h exactly matching the number of trading days covered by the associated implied volatility, but for notational convenience we continue to write $h = 22$ for all monthly realized measures. Finally, IV_t denotes the implied volatility backed out from the price of

Table 1: Realized volatility HAR models

Const.	$RV_{t-22,t}$	$RV_{t-5,t}$	RV_t	$C_{t-22,t}$	$C_{t-5,t}$	C_t	$J_{t-22,t}$	$J_{t-5,t}$	J_t	IV_t	Adj R^2	AR_{12}	MAFE
Panel A: Foreign exchange data													
0.0061	0.2186	0.0981	0.1706	–	–	–	–	–	–	–	26.0%	9.92	0.3376
(0.0011)	(0.1138)	(0.1438)	(0.0828)										
0.0061	–	–	–	0.2355	0.0871	0.2407	0.1922	−0.8014	0.0898	–	26.9%	12.05	0.3774
(0.0011)				(0.1597)	(0.1623)	(0.0930)	(0.6783)	(0.5640)	(0.2193)				
0.0022	–	–	–	–	–	–	–	–	–	0.8917	40.7%	15.64	0.2938
(0.0011)										(0.0884)			
0.0021	−0.1483	0.0769	0.0765	–	–	–	–	–	–	0.8733	41.1%	21.50^{*}	0.3124
(0.0011)	(0.1178)	(0.1284)	(0.0754)							(0.1419)			
0.0022	–	–	–	−0.0517	0.0097	0.1114	−0.7076	0.0996	−0.0474	0.8715	40.4%	23.90^{*}	0.3496
(0.0012)				(0.1526)	(0.1471)	(0.0869)	(0.6319)	(0.5326)	(0.1993)	(0.1515)			
Panel B: S&P 500 data													
Const.	$RV_{t-22,t}$	$RV_{t-5,t}$	RV_t	$C_{t-22,t}$	$C_{t-5,t}$	C_t	$J_{t-22,t}$	$J_{t-5,t}$	J_t	IV_t	Adj R^2	AR_{12}	MAFE
0.0053	0.6240	−0.3340	0.6765	–	–	–	–	–	–	–	53.0%	41.50^{**}	3.2011
(0.0025)	(0.1132)	(0.1039)	(0.1007)										
0.0037	–	–	–	0.1568	0.0407	0.9646	0.2727	−0.1427	−0.7903	–	61.9%	25.08^{*}	2.8250
(0.0023)				(0.1327)	(0.1353)	(0.1088)	(0.2720)	(0.1812)	(0.3873)				
−0.0050	–	–	–	–	–	–	–	–	–	1.0585	62.1%	17.83	1.9912
(0.0027)										(0.0667)			
−0.0052	0.0378	−0.1617	0.3177	–	–	–	–	–	–	0.9513	64.0%	26.18^{*}	2.5089
(0.0027)	(0.1311)	(0.0943)	(0.1026)							(0.1391)			
−0.0051	–	–	–	−0.1511	0.0633	0.6016	−0.4486	0.0454	−0.7019	0.7952	68.2%	24.87^{*}	2.2582
(0.0027)				(0.1336)	(0.1237)	(0.1194)	(0.2809)	(0.1690)	(0.3541)	(0.1447)			
Panel C: Treasury bond data													
Const.	$RV_{t-22,t}$	$RV_{t-5,t}$	RV_t	$C_{t-22,t}$	$C_{t-5,t}$	C_t	$J_{t-22,t}$	$J_{t-5,t}$	J_t	IV_t	Adj R^2	AR_{12}	MAFE
0.0031	0.3600	0.1112	0.1389	–	–	–	–	–	–	–	32.5%	20.30	0.2253
(0.0005)	(0.1106)	(0.1143)	(0.0744)										
0.0037	–	–	–	0.4203	0.1436	0.0826	−0.1502	−0.3379	0.3660	–	37.0%	19.17	0.2253
(0.0005)				(0.1347)	(0.1363)	(0.0776)	(0.2792)	(0.2567)	(0.1764)				
0.0023	–	–	–	–	–	–	–	–	–	0.5686	35.0%	29.45^{**}	0.2023
(0.0006)										(0.0641)			
0.0018	0.0462	0.1835	0.0817	–	–	–	–	–	–	0.3933	40.4%	22.75^{*}	0.2092
(0.0006)	(0.1254)	(0.1086)	(0.0710)							(0.0882)			
0.0023	–	–	–	0.1736	0.1424	0.0318	−0.7019	−0.0605	0.2812	0.4129	45.5%	21.70^{*}	0.2149
(0.0006)				(0.1355)	(0.1267)	(0.0729)	(0.2842)	(0.2457)	(0.1649)	(0.0867)			

Note: The table shows HAR-RV-CJIV results for the specification (14) with standard errors in parentheses. Adj R^2 denotes the adjusted R^2 for the regression and AR_{12} is the LM test statistic (with 12 lags) for the null of no serial correlation in the residuals. One and two asterisks denote rejection at the 5% and 1% significance levels, respectively. The last column (MAFE) reports out-of-sample mean absolute forecast errors ($\times 100$) for 24 rolling window one-step ahead forecasts based on $n-24$ observations.

the relevant one-month option sampled on day t, and is in variance form.

The HAR-RV-CJIV model is the Mincer & Zarnowitz (1969) type regression

$$RV_{t,t+22} = \alpha + \gamma_m x_{t-22,t} + \gamma_w x_{t-5,t} + \gamma_d x_t + \beta IV_t + \varepsilon_{t,t+22}, \quad t = 22, 44, 66, \ldots, 22n, \tag{14}$$

where $\varepsilon_{t,t+22}$ is the monthly forecasting error, and $x_{t-h,t}$ is either $RV_{t-h,t}$ or the vector $(C_{t-h,t}, J_{t-h,t})$. When a variable is not included in the specific regression, $\beta = 0$ or $\gamma_m = \gamma_w = \gamma_d = 0$ is imposed. Note that $x_{t-22,t}$ contains lagged realized volatility measures covering a month, whereas $x_{t-5,t}$ and x_t allow extracting information about future volatility from the more recent (one week and one day) history of past squared returns.

Table 1 shows the results for the HAR-RV-CJIV model in (14). We report coefficient estimates (standard errors in parentheses) together with adjusted R^2 and the Breusch-Godfrey LM test for residual autocorrelation up to lag 12 (one year), denoted AR_{12}, used here instead of the standard Durbin-Watson statistic due to the presence of lagged endogenous variables in several of the regressions. Under the null hypothesis of absence of serial dependence in the residuals, the AR_{12} statistic is asymptotically χ^2 with 12 degrees of freedom. The last column (MAFE) reports out-of-sample mean absolute forecast errors

($\times 100$) for 24 rolling window one-step ahead forecasts based on $n-24$ observations.

In the first line of each panel, $x = RV$, so this is a monthly frequency HAR-RV model (Corsi (2009)). In the foreign exchange market (Panel A), the combined impact from the monthly, weekly, and daily RV on future realized volatility is $.22+.10+.17 = .49$, strikingly close to the first order autocorrelation of monthly RV, which is .46. The t-statistics are 1.92, .68, and 2.06, respectively, indicating that the weekly variable contains only minor information concerning future monthly exchange rate volatility. In the stock market (Panel B), all three RV measures are significant, but the weekly measure with a negative coefficient. In contrast to the foreign exchange market, the AR_{12} test is significant. Panel C is for bond data and the results in the first line show that only monthly RV is significant.

In row two of each panel of Table 1, $x = (C, J)$, so this is a monthly frequency HAR-RV-CJ model (Andersen et al. (2007)). The conclusions for C are similar to those for RV in the first row, except that the monthly and weekly components become insignificant in the stock market. The jump components are insignificant, except daily J in the stock and bond markets. Adjusted R^2 improves when moving from first to second line of each panel of Table 1, thus confirming the enhanced in-sample (Mincer-Zarnowitz) forecasting performance obtained by splitting RV into its separate components also found by Andersen et al. (2007). Out-of-sample forecasting performance (MAFE) improves in the stock market when separating C and J, but remains unchanged in the bond market, and actually deteriorates in the currency market, hence showing the relevance of including this criterion in the analysis.

Next, implied volatility is added to the information set at time t in the HAR regressions. When RV is included together with IV, fourth row, all the realized volatility coefficients turn insignificant in the foreign exchange and bond markets, and only daily RV remains significant in the stock market. Indeed, IV gets t-statistics of 6.15, 6.84, and 4.46 in the three markets, providing clear evidence of the relevance of IV in forecasting future volatility. The last row of each panel shows the results when including C and J together with IV, i.e., the full HAR-RV-CJIV model (14). In the foreign exchange market, IV completely subsumes the information content of both C and J at all frequencies. Adjusted R^2 is about equally high in the third line of the panel, where IV is the sole regressor and where also MAFE takes the best (lowest) value in the panel. In the stock market, both daily components of RV remain significant, and the adjusted R^2 increases from 62% to 68% relative to having IV as the sole regressor, but again MAFE points to the specification with only IV included as the best forecast. In the bond market, IV gets the highest t-statistic, as in the other two markets. In this case, the monthly jump component $J_{t-22,t}$ is also significant and adjusted R^2 improves markedly, both between lines three and four (adding realized volatility) and between lines four and five (separating the RV components), but the ordering by MAFE is the reverse. The AR_{12} test does show mild signs of misspecification in all three markets.

The finding so far is that IV as a forecast of future volatility contains incremental information relative to return-based forecasts in all three markets, even when using the new nonparametric jump separation technique for RV and exploiting potential forecasting power of the C and J components at different frequencies using the HAR methodology.

Table 2: Continuous component HAR models

Const.	$C_{t-22,t}$	$C_{t-5,t}$	C_t	$J_{t-22,t}$	$J_{t-5,t}$	J_t	IV_t	Adj R^2	AR_{12}	MAFE
Panel A: Foreign exchange data										
0.0053	0.2273	0.0634	0.1892	–	–	–	–	28.6%	9.23	0.2987
(0.0009)	(0.1128)	(0.1370)	(0.0779)							
0.0053	0.2651	0.0487	0.2176	−0.1025	−0.5571	0.0826	–	27.9%	10.98	0.3238
(0.0009)	(0.1410)	(0.1433)	(0.0821)	(0.5990)	(0.4980)	(0.1936)				
0.0020	–	–	–	–	–	–	0.7857	39.9%	18.84	0.2632
(0.0010)							(0.0791)			
0.0020	0.0209	−0.0171	0.1077	−0.8674	0.2088	−0.0340	0.7409	40.2%	24.24*	0.3129
(0.0010)	(0.1359)	(0.1310)	(0.0774)	(0.5628)	(0.4743)	(0.1775)	(0.1349)			
Panel B: S&P 500 data										
0.0031	0.1556	0.0959	0.7356	–	–	–	–	67.2%	25.39*	2.2421
(0.0018)	(0.0989)	(0.0967)	(0.0783)							
0.0023	0.1246	0.0853	0.8636	0.3412	−0.3477	−0.4236	–	73.7%	14.25	1.8550
(0.0017)	(0.0944)	(0.0962)	(0.0774)	(0.1934)	(0.1288)	(0.2753)				
−0.0052	–	–	–	–	–	–	0.9498	68.5%	24.99*	1.4481
(0.0021)							(0.0520)			
−0.0040	−0.0945	0.1014	0.6053	−0.1721	−0.2138	−0.3606	0.5659	78.1%	20.23	1.5407
(0.0019)	(0.0949)	(0.0879)	(0.0848)	(0.1997)	(0.1201)	(0.2517)	(0.1029)			
Panel C: Treasury bond data										
0.0023	0.4028	0.1480	0.0959	–	–	–	–	39.4%	27.14**	0.2251
(0.0004)	(0.1216)	(0.1236)	(0.0722)							
0.0030	0.4220	0.1082	0.1171	−0.4624	−0.3924	0.2856	–	43.1%	22.38*	0.2199
(0.0005)	(0.1221)	(0.1236)	(0.0703)	(0.2531)	(0.2327)	(0.1599)				
0.0018	–	–	–	–	–	–	0.4977	29.4%	42.16**	0.2145
(0.0006)							(0.0637)			
0.0019	0.2353	0.1073	0.0786	−0.8798	−0.1826	0.2214	0.3124	48.3%	20.50	0.2065
(0.0005)	(0.1259)	(0.1177)	(0.0678)	(0.2641)	(0.2282)	(0.1532)	(0.0806)			

Note: The table shows HAR-C-CJIV results for (15), using the same definitions and layout as in Table 1.

Indeed, in the preferred specification for the foreign exchange market, all realized measures are insignificant or left out, showing that the conclusions of Jorion (1995), Xu & Taylor (1995), and Pong et al. (2004) hold up even to these new improvements of the return-based forecasts. In fact, the superiority of the implied volatility forecast of future volatility extends to all three markets based on out-of-sample forecasting, where MAFE within each panel in Table 1 is minimized by the specification with IV as the sole forecasting variable.

5.2 Forecasting the Continuous and Jump Components

We now split $RV_{t,t+22}$ on the left hand side of (14) into its continuous and jump components, $C_{t,t+22}$ and $J_{t,t+22}$, and forecast these separately. This is particularly interesting because Andersen et al. (2007) show that the two components exhibit very different time series properties, as also evident from Figure 1. Thus, forecasting should be carried out in different manners for the two, and we modify the HAR methodology accordingly. Andersen et al. (2007) did not consider separate forecasting of the components. Since our specifications include implied volatility as well, we are able to assess the incremental information in option prices on the future continuous and jump components of volatility separately.

Our HAR-C-CJIV model for forecasting the future continuous component is

$$C_{t,t+22} = \alpha + \gamma_m x_{t-22,t} + \gamma_w x_{t-5,t} + \gamma_d x_t + \beta IV_t + \varepsilon_{t,t+22}, \quad t = 22, 44, 66, \ldots, 22n, \quad (15)$$

where $C_{t,t+22}$ replaces $RV_{t,t+22}$ as regressand compared to (14) and x now contains either C or the vector (C, J). Table 2 shows the results in the same format as in Table 1.

The results in Table 2 are similar to those in Table 1. This confirms that C is as amenable to forecasting as RV, hence demonstrating the value of the new approach of separate HAR modeling of C and J. The AR_{12} tests show only modest signs of misspecification, except in the bond market. In the foreign exchange and bond markets, adjusted R^2s are similar to comparable specifications in Table 1, and in the stock market they are higher than in Table 1. Implied volatility generally gets higher coefficients and t-statistics than the lagged C and J components, and adjusted R^2 is highest when IV is included along with these. In the foreign exchange market, IV completely subsumes the information content of the realized measures, just as in Panel A of Table 1. The out-of-sample forecasting evidence based on MAFE suggests using IV as the sole forecasting variable in the foreign exchange and stock markets, and combining IV with the realized measures in the bond market.

We next consider the predictability of the jump component of realized volatility. The relevant HAR-J-CJIV model takes the form

$$J_{t,t+22} = \alpha + \gamma_m x_{t-22,t} + \gamma_w x_{t-5,t} + \gamma_d x_t + \beta IV_t + \varepsilon_{t,t+22}, \quad t = 22, 44, 66, \ldots, 22n, \quad (16)$$

where x now contains either J or (C, J). Table 3 reports the results. Specifically, in line three of each panel, IV is used to predict the jump component of future volatility. It is strongly significant in all three markets and gets higher t-statistics than all other variables considered. The highest adjusted R^2s in the table are obtained in the fourth line of each panel, where all variables are included. Here, the AR_{12} test shows no signs of misspecification in the foreign exchange and bond markets, although it rejects in the stock market. Implied volatility remains highly significant in all three markets and turns out to be the strongest predictor of future jumps (in terms of t-statistics) even when the C and J components at all frequencies are included. The coefficient on IV ranges between .07 and .23 across markets and specifications, consistent with the mean jump component being an order of magnitude smaller than implied volatility in Figure 1. Indeed, in the bond market, IV subsumes the information content of both C and J at all frequencies.

In the foreign exchange and stock markets, the out-of-sample forecasting evidence suggests that IV should be used alone as the sole forecasting variable even when forecasting only the future jump component. In the bond market, the MAFE criterion selects the forecast using only past jump components.

Comparing across Tables 1-3, the results are most similar in Tables 1 and 2, so RV and C behave similarly also in this forecasting context, and our results show that IV is important in forecasting both. The difference in results when moving to Table 3 reinforces that C and J should be treated separately. When doing so, we find that, firstly, jumps are predictable from variables in the information set, and, secondly, IV retains incremental information, thus suggesting that option prices incorporate jump information. Finally, the out-of-sample forecasting evidence suggests using IV as the sole forecasting variable in the foreign exchange and stock markets, whether forecasting RV or either of its components.

5.3 The Vector Heterogeneous Autoregressive (VecHAR) Model

We now introduce a simultaneous system approach for joint analysis of IV, C, and J. The reason is that, firstly, the results up to this point have been obtained in different regression equations that are not independent, and some relevant joint hypotheses involve

Table 3: Jump component HAR models

Const.	$C_{t-22,t}$	$C_{t-5,t}$	C_t	$J_{t-22,t}$	$J_{t-5,t}$	J_t	IV_t	Adj R^2	AR_{12}	MAFE
Panel A: Foreign exchange data										
0.0011	–	–	–	0.3272	−0.1192	−0.0167	–	7.3%	14.60	0.0607
(0.0001)				(0.0908)	(0.0910)	(0.0372)				
0.0009	−0.0295	0.0384	0.0231	0.2947	−0.2443	0.0071	–	14.4%	15.15	0.0616
(0.0002)	(0.0265)	(0.0270)	(0.0155)	(0.1127)	(0.0937)	(0.0364)				
0.0002	–	–	–	–	–	–	0.1060	24.1%	10.83	0.0529
(0.0002)							(0.0154)			
0.0003	−0.0726	0.0268	0.0038	0.1598	−0.1092	−0.0134	0.1307	27.2%	11.09	0.0539
(0.0002)	(0.0259)	(0.0250)	(0.0148)	(0.1073)	(0.0904)	(0.0338)	(0.0257)			
Panel B: S&P 500 data										
Const.	$C_{t-22,t}$	$C_{t-5,t}$	C_t	$J_{t-22,t}$	$J_{t-5,t}$	J_t	IV_t	Adj R^2	AR_{12}	MAFE
0.0031	–	–	–	0.0041	0.2036	−0.3052	–	4.0%	63.65**	1.1983
(0.0008)				(0.0992)	(0.0770)	(0.1571)				
0.0015	0.0321	−0.0446	0.1010	−0.0685	0.2049	−0.3668	–	6.2%	57.36**	1.1494
(0.0010)	(0.0577)	(0.0588)	(0.0473)	(0.1182)	(0.0788)	(0.1684)				
0.0002	–	–	–	–	–	–	0.1087	8.0%	59.37**	1.0042
(0.0012)							(0.0288)			
−0.0011	−0.0567	−0.0381	−0.0037	−0.2765	0.2592	−0.3413	0.2293	12.7%	51.56**	1.1276
(0.0012)	(0.0613)	(0.0568)	(0.0548)	(0.1290)	(0.0776)	(0.1626)	(0.0664)			
Panel C: Treasury bond data										
Const.	$C_{t-22,t}$	$C_{t-5,t}$	C_t	$J_{t-22,t}$	$J_{t-5,t}$	J_t	IV_t	Adj R^2	AR_{12}	MAFE
0.0007	–	–	–	0.3100	0.0438	0.0820	–	12.2%	17.12	0.0463
(0.0001)				(0.0867)	(0.0808)	(0.0553)				
0.0007	−0.0017	0.0354	−0.0345	0.3122	0.0546	0.0804	–	11.6%	21.00	0.0516
(0.0002)	(0.0430)	(0.0436)	(0.0248)	(0.0892)	(0.0821)	(0.0564)				
0.0005	–	–	–	–	–	–	0.0709	6.9%	34.70**	0.0640
(0.0002)							(0.0207)			
0.0004	−0.0617	0.0351	−0.0469	0.1780	0.1221	0.0598	0.1005	18.2%	19.46	0.0534
(0.0002)	(0.0448)	(0.0419)	(0.0241)	(0.0940)	(0.0812)	(0.0545)	(0.0287)			

Note: The table shows HAR-J-CJIV results for (16), using the same definitions and layout as in Table 1.

cross-equation restrictions. Secondly, IV may be measured with error stemming from non-synchronous option and futures prices, misspecification of the option pricing formula, etc. Such errors-in-variable problems generate correlation between regressor and error terms in the forecasting equations for C and J, and thus an endogeneity problem. In addition, the realized measures contain sampling error, as discussed in section 2. Our simultaneous system approach provides an efficient method for handling these endogeneity issues.[3]

Thus, we consider the vector heterogeneous autoregressive (VecHAR) system

$$\begin{pmatrix} 1 & 0 & -\beta_1 \\ 0 & 1 & -\beta_2 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} C_{t,t+22} \\ J_{t,t+22} \\ IV_t \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \end{pmatrix} + \begin{pmatrix} A_{11m} & A_{12m} & 0 \\ A_{21m} & A_{22m} & 0 \\ A_{31m} & A_{32m} & A_{33m} \end{pmatrix} \begin{pmatrix} C_{t-22,t} \\ J_{t-22,t} \\ IV_{t-1} \end{pmatrix} \tag{17}$$

$$+ \begin{pmatrix} A_{11w} & A_{12w} \\ A_{21w} & A_{22w} \\ A_{31w} & A_{32w} \end{pmatrix} \begin{pmatrix} C_{t-5,t} \\ J_{t-5,t} \end{pmatrix} + \begin{pmatrix} A_{11d} & A_{12d} \\ A_{21d} & A_{22d} \\ A_{31d} & A_{32d} \end{pmatrix} \begin{pmatrix} C_t \\ J_t \end{pmatrix} + \begin{pmatrix} \varepsilon^1_{t,t+22} \\ \varepsilon^2_{t,t+22} \\ \varepsilon^3_{t,t+22} \end{pmatrix}.$$

The first two equations comprise the forecasting equations (15) and (16) for C and J, and the third endogenizes IV. There are two sources of simultaneity in the VecHAR system. Firstly, the off-diagonal terms β_1 and β_2 in the leading coefficient matrix accommodates dependence of $C_{t,t+22}$ and $J_{t,t+22}$ on the endogenous variable IV_t. Secondly, the system errors may be contemporaneously correlated. In the third equation, option prices may reflect return movements during the previous month, and via the HAR type specification more re-

[3]Engle & Gallo (2006) also consider a trivariate system, but for three realized quadratic variation measures, and without the HAR feature or jumps. Implied volatility is also not included in their system, although it is in subsequent univariate regressions.

Table 4: VecHAR models

Panel A: Foreign exchange data											
Dep. var.	Constant	$C_{t-22,t}$	$C_{t-5,t}$	C_t	$J_{t-22,t}$	$J_{t-5,t}$	J_t	IV_t	IV_{t-1}	AR_{12}	MAFE
$C_{t,t+22}$	0.0024	0.0552	−0.0078	0.1231	−0.7601	0.1013	−0.0177	0.6369	–	33.88^{**}	0.3114
	(0.0010)	(0.1770)	(0.1418)	(0.0721)	(0.5914)	(0.6028)	(0.1441)	(0.1886)			
$J_{t,t+22}$	0.0004	−0.0658	0.0286	0.0068	0.1811	−0.1305	−0.0102	0.1100	–	15.16	0.0550
	(0.0002)	(0.0356)	(0.0311)	(0.0139)	(0.1191)	(0.1037)	(0.0258)	(0.0383)			
IV_t	0.0013	−0.0668	−0.0521	0.0788	−0.0431	−0.3046	0.0242	–	0.9028	8.57	0.0766
	(0.0006)	(0.0547)	(0.0405)	(0.0365)	(0.1431)	(0.1874)	(0.0436)		(0.0665)		
Panel B: S&P 500 data											
Dep. var.	Constant	$C_{t-22,t}$	$C_{t-5,t}$	C_t	$J_{t-22,t}$	$J_{t-5,t}$	J_t	IV_t	IV_{t-1}	AR_{12}	MAFE
$C_{t,t+22}$	−0.0054	−0.1432	0.1050	0.5478	−0.2863	−0.1840	−0.3466	0.6918	–	20.49	1.5277
	(0.0022)	(0.1307)	(0.0882)	(0.1658)	(0.1902)	(0.1647)	(0.3820)	(0.1599)			
$J_{t,t+22}$	−0.0015	−0.0732	−0.0369	−0.0232	−0.3152	0.2693	−0.3365	0.2720	–	59.32^{**}	1.1225
	(0.0012)	(0.0980)	(0.0429)	(0.0429)	(0.1474)	(0.1895)	(0.3471)	(0.1018)			
IV_t	0.0009	−0.0144	0.0159	0.0307	0.3800	−0.0158	−0.3868	–	0.9349	28.36^{**}	0.6819
	(0.0007)	(0.0613)	(0.0454)	(0.0358)	(0.0547)	(0.0450)	(0.0956)		(0.0553)		
Panel C: Treasury bond data											
Dep. var.	Constant	$C_{t-22,t}$	$C_{t-5,t}$	C_t	$J_{t-22,t}$	$J_{t-5,t}$	J_t	IV_t	IV_{t-1}	AR_{12}	MAFE
$C_{t,t+22}$	0.0019	0.2253	0.1073	0.0766	−0.9023	−0.1713	0.2180	0.3292	–	20.05	0.2042
	(0.0005)	(0.1209)	(0.1089)	(0.0646)	(0.2598)	(0.1815)	(0.0609)	(0.0806)			
$J_{t,t+22}$	0.0003	−0.0740	0.0350	−0.0494	0.1506	0.1358	0.0556	0.1210	–	17.60	0.0539
	(0.0002)	(0.0479)	(0.0412)	(0.0289)	(0.1142)	(0.0765)	(0.0589)	(0.0525)			
IV_t	0.0001	0.0328	−0.0264	0.0517	0.2547	−0.0344	0.0448	–	0.9172	21.85^{*}	0.0404
	(0.0003)	(0.0447)	(0.0375)	(0.0355)	(0.2282)	(0.1091)	(0.0389)		(0.0400)		

Note: The table shows FIML results for the simultaneous VecHAR system (17) with robust (sandwich-formula) standard errors in parentheses. AR_{12} and MAFE are defined as in Table 1.

cent returns may receive higher weight. In addition, our specification allows dependence on IV_{t-1}, i.e., one-day lagged implied volatility sampled on Monday for the same option contract as in IV_t, which is sampled on Tuesday. The specification of the third equation is similar to using IV_{t-1} as an additional instrument for IV_t in an instrumental variables treatment of the endogeneity problem, but the system approach in (17) is more general and efficient.

In Table 4 we present the results of Gaussian full information maximum likelihood (FIML) estimation of the VecHAR system with robust standard errors (sandwich-formula, $H^{-1}VH^{-1}$, where H is the Hessian and V the outer-product-gradient matrix) in parentheses. Of course, the results are asymptotically valid even in the absence of Gaussianity. The AR_{12} tests show only mild signs of misspecification in the foreign exchange and bond markets, although the tests are significant in two of the equations for the stock market.

Implied volatility is strongly significant in the forecasting equations for both C and J in all three markets, showing that option prices contain incremental information beyond that in high-frequency realized measures. In the foreign exchange market, IV subsumes the information content of all other variables in forecasting both C and J.

Out-of-sample forecasting performance actually improves for C (first equation) in the VecHAR system relative to comparable univariate specifications (last row of each panel in Table 2). For J, out-of-sample forecasting performance is similar in the VecHAR system and in the last row of each panel in Table 3, with a small improvement in the simultaneous system for the stock market and a small deterioration in the other two markets.

Table 5 shows results of likelihood ratio (LR) tests of many hypotheses of interest in the VecHAR model. First, the hypothesis $\mathrm{H}_2 : A_{11w} = 0, A_{12w} = 0$ in (17) is the relevant forecasting efficiency hypothesis in the C equation with respect to both weekly realized components. This hypothesis is not rejected in any of the markets. Indeed, in the foreign exchange market, IV subsumes the information content of C and J at all frequencies, with

Table 5: LR tests in VecHAR models

	Panel A: Currency data			Panel B: S&P 500 data			Panel C: Bond data		
Hypothesis	LR	d.f.	p-values	LR	d.f.	p-values	LR	d.f.	p-values
$\mathrm{H}_1 : A_{11m} = 0, A_{12m} = 0$	1.8922	2	0.3883	2.9715	2	0.2263	14.564	2	0.0007
$\mathrm{H}_2 : A_{11w} = 0, A_{12w} = 0$	0.0490	2	0.9758	3.8762	2	0.1440	1.4315	2	0.4888
$\mathrm{H}_3 : A_{11d} = 0, A_{12d} = 0$	2.6273	2	0.2688	32.824	2	0.0000	3.2478	2	0.1971
$\mathrm{H}_4 : \beta_1 = 1$	5.8950	1	0.0152	6.7789	1	0.0092	46.199	1	0.0000
$\mathrm{H}_5 : A_{11m} = 0, A_{12m} = 0, \beta_1 = 1$	14.560	3	0.0022	20.427	3	0.0001	85.798	3	0.0000
$\mathrm{H}_6 : A_{11w} = 0, A_{12w} = 0, \beta_1 = 1$	7.1047	3	0.0686	9.1775	3	0.0270	47.871	3	0.0000
$\mathrm{H}_7 : A_{11d} = 0, A_{12d} = 0, \beta_1 = 1$	6.9558	3	0.0733	35.498	3	0.0000	46.400	3	0.0000
$\mathrm{H}_8 : \bar{\mathbf{A}}_m = 0, \bar{\mathbf{A}}_w = 0, \bar{\mathbf{A}}_d = 0$	29.377	12	0.0035	81.304	12	0.0000	73.212	12	0.0000
$\mathrm{H}_9 : \beta_2 = 0$	13.227	1	0.0003	13.403	1	0.0003	15.036	1	0.0001
$\mathrm{H}_{10} : \beta_1 + \beta_2 = 1$	2.2915	1	0.1301	0.0512	1	0.8210	29.054	1	0.0000

Note: The table shows LR test results for the simultaneous VecHAR system (17) where the matrix notation $\bar{\mathbf{A}}_k = \begin{pmatrix} A_{11k} & A_{12k} \\ A_{21k} & A_{22k} \end{pmatrix}$, $k = m, w, d$, is used.

p-values of $\mathrm{H}_1 : A_{11m} = 0, A_{12m} = 0$ and $\mathrm{H}_3 : A_{11d} = 0, A_{12d} = 0$ of 39% and 27% in this market. IV subsumes the information content of the monthly measures (H_1) in the stock market, and the daily measures (H_3) in the bond market. Unbiasedness of IV as a forecast of C ($\mathrm{H}_4 : \beta_1 = 1$) is rejected at the 5% level in all three markets. The estimated coefficient on IV is below unity in all three markets in Table 4, showing that IV is upward biased as a forecast of future C. Possible reasons for this phenomenon are that volatility risk is priced (c.f. Bollerslev & Zhou (2006)) or that IV reflects information about future J as well, which we return to in H_9 and H_{10}.

In H_5-H_7, the unbiasedness hypothesis H_4 is tested jointly with the efficiency hypotheses H_1-H_3. Consistent with previous results, H_5-H_7 are strongly rejected in the stock and bond markets and H_6-H_7 (efficiency with respect to daily and weekly measures along with unbiasedness) are not rejected at the 5% level in the foreign exchange market.

Next, we consider cross-equation restrictions which hence require the system approach. Using the matrix notation $\bar{\mathbf{A}}_k = \begin{pmatrix} A_{11k} & A_{12k} \\ A_{21k} & A_{22k} \end{pmatrix}$, $k = m, w, d$, we examine in $\mathrm{H}_8 : \bar{\mathbf{A}}_m = 0, \bar{\mathbf{A}}_w = 0, \bar{\mathbf{A}}_d = 0$ the hypothesis that all realized components in both the continuous and jump equations are jointly insignificant. This is rejected in all three markets.

In $\mathrm{H}_9 : \beta_2 = 0$, we examine the hypothesis that IV carries no incremental information about future J, relative to the realized measures. This is strongly rejected in all three markets. Finally, in $\mathrm{H}_{10} : \beta_1 + \beta_2 = 1$, again a cross-equation restriction, we test the hypothesis that IV_t is an unbiased forecast of total realized volatility, $RV_{t,t+22} = C_{t,t+22} + J_{t,t+22}$. Although unbiasedness of IV as a forecast of future C, $\mathrm{H}_4 : \beta_1 = 1$, is rejected at the 5% level or better in all markets, $\mathrm{H}_{10} : \beta_1 + \beta_2 = 1$ is not rejected in the foreign exchange and stock markets. This reinforces earlier conclusions that IV does forecast more than just the continuous component, that jumps are, to some extent, predictable, and, indeed, that option prices are calibrated to incorporate information about future jumps.

6 Conclusions and Discussion

This paper examines the role of implied volatility in forecasting future realized volatility and jumps in the foreign exchange, stock, and bond markets. Realized volatility is separated into its continuous sample path and jump components, since Andersen et al.

(2007) show that this leads to improved forecasting performance. We assess the incremental forecasting power of implied volatility relative to Andersen et al. (2007).

On the methodological side, we apply the HAR model proposed by Corsi (2009) and applied by Andersen et al. (2007). We include implied volatility as an additional regressor, and also consider forecasting of the separate continuous and jump components of realized volatility. Furthermore, we introduce a vector HAR (VecHAR) model for simultaneous modeling of implied volatility and the separate components of realized volatility, controlling for possible endogeneity issues.

On the substantive side, our empirical results using both in-sample Mincer & Zarnowitz (1969) regressions and out-of-sample forecasting show that in all three markets, option implied volatility contains incremental information about future return volatility relative to both the continuous and jump components of realized volatility. Indeed, implied volatility subsumes the information content of several realized measures in all three markets. In addition, implied volatility is an unbiased forecast of the sum of the continuous and jump components, i.e., of total realized volatility, in the foreign exchange and stock markets. The out-of-sample forecasting evidence confirms that implied volatility should be used in forecasting future realized volatility or the continuous component of this in all three markets. Finally, our results show that even the jump component of realized return volatility is, to some extent, predictable, and that option implied volatility enters significantly in the relevant jump forecasting equation for all three markets.

Overall, our results are interesting and complement the burgeoning realized volatility literature. What we show is that implied volatility generally contains additional ex-ante information on volatility and its continuous sample path and jump components beyond that in realized volatility and its components. This ex-ante criterion is not everything that realized volatility may be used for, and it is possibly not the most important use. For example, realized volatility and its components can be used for ex-post assessments of what volatility has been, whether there have been jumps in prices or not, etc. Implied volatility (even ex-post implied volatility) is not well suited for these purposes.

Acknowledgements

We are grateful to the editors, two anonymous referees, Torben G. Andersen, John Y. Campbell, Per Frederiksen, Pierre Perron, Jeremy Stein, Matti Suominen, Luis Viceira, and participants at various conferences and seminars for useful comments, and especially to Tim Bollerslev for extensive comments and for providing the realized volatility, bipower variation, and tripower quarticity data. We thank Christian Bach, Rolf Karlsen, and Christian Sønderup for research assistance and the Center for Analytical Economics (CAE) at Cornell, the Center for Analytical Finance (CAF) at Aarhus, the Center for Research in Econometric Analysis of Time Series (CREATES, funded by the Danish National Research Foundation) at Aarhus, the Social Sciences and Humanities Research Council of Canada (SSHRC), and the Danish Social Science Research Council (FSE) for research support. This paper combines and further develops the material in our previous three papers Chris-

tensen & Nielsen (2005) and Busch, Christensen & Nielsen (2005, 2006).

References

Amin, K.I., Morton, A.J., 1994. Implied volatility functions in arbitrage-free term structure models. Journal of Financial Economics 35, 141–180.

Andersen, T.G., Bollerslev, T., Diebold, F.X., 2007. Roughing it up: including jump components in the measurement, modeling, and forecasting of return volatility. Review of Economics and Statistics 89, 701–720.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2003. Modeling and forecasting realized volatility. Econometrica 71, 579–625.

Andersen, T.G., Bollerslev, T., Meddahi, N., 2004. Analytical evaluation of volatility forecasts. International Economic Review 45, 1079–1110.

Andersen, T.G., Bondarenko, O., 2007. Construction and interpretation of model-free implied volatility, in I. Nelken, ed., Volatility as an Asset Class. Risk Books, London, UK, pp. 141–181.

Barndorff-Nielsen, O.E., Graversen, S.E., Jacod, J., Shephard, N., 2006. Limit theorems for realised bipower variation in econometrics. Econometric Theory 22, 677–719.

Barndorff-Nielsen, O.E., Shephard, N., 2004. Power and bipower variation with stochastic volatility and jumps (with discussion). Journal of Financial Econometrics 2, 1–57.

Barndorff-Nielsen, O.E., Shephard, N., 2006. Econometrics of testing for jumps in financial economics using bipower variation. Journal of Financial Econometrics 4, 1–30.

Barndorff-Nielsen, O.E., Shephard, N., 2007. Variation, jumps, market frictions, and high frequency data in financial econometrics, in R. Blundell, T. Persson, W.K. Newey, eds., Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress. Cambridge University Press, Cambridge, UK, pp. 328–372.

Barndorff-Nielsen, O.E., Shephard, N., Winkel, M., 2006. Limit theorems for multipower variation in the presence of jumps. Stochastic Processes and Their Applications 116, 796–806.

Bates, D.S., 1996a. Dollar jump fears, 1984-1992: distributional abnormalities implicit in currency futures options. Journal of International Money and Finance 15, 65–93.

Bates, D.S., 1996b. Jumps and stochastic volatility: exchange rate processes implicit in Deutsche Mark options. Review of Financial Studies 9, 69–107.

Black, F., 1976. The pricing of commodity contracts. Journal of Financial Economics 3, 167–179.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–654.

Blair, B.J., Poon, S., Taylor, S.J., 2001. Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns. Journal of Econometrics 105, 5–26.

Bollerslev, T., Zhou, H., 2006. Volatility puzzles: a simple framework for gauging return-volatility regressions. Journal of Econometrics 131, 123–150.

Busch, T., Christensen, B.J., Nielsen, M.Ø., 2005. Forecasting exchange rate volatility in the presence of jumps. QED Working Paper 1187, Queen's University.

Busch, T., Christensen, B.J., Nielsen, M.Ø., 2006. The information content of Treasury bond options concerning future volatility and price jumps. QED Working Paper 1188,

Queen's University.
Canina, L., Figlewski, S., 1993. The informational content of implied volatility. Review of Financial Studies 6, 659–681.
Christensen, B.J., Nielsen, M.Ø., 2005. The implied-realized volatility relation with jumps in underlying asset prices. QED Working Paper 1186, Queen's University.
Christensen, B.J., Prabhala, N.R., 1998. The relation between implied and realized volatility. Journal of Financial Economics 50, 125–150.
Corsi, F., 2009. A simple approximate long memory model of realized volatility. Journal of Financial Econometrics 7, 174–196.
Covrig, V., Low, B.S., 2003. The quality of volatility traded on the over-the-counter market: a multiple horizons study. Journal of Futures Markets 23, 261–285.
Cox, J.C., Ingersoll, J.E., Ross, S.A., 1981. The relationship between forward prices and futures prices. Journal of Financial Economics 9, 321–346.
Day, T.E., Lewis, C.M., 1992. Stock market volatility and the information content of stock index options. Journal of Econometrics 52, 267–287.
Engle, R.F., Gallo, G.M., 2006. A multiple indicators model for volatility using intra-daily data. Journal of Econometrics 131, 3–27.
Fleming, J., 1998. The quality of market volatility forecasts implied by S&P 100 index option prices. Journal of Empirical Finance 5, 317–345.
French, D.W., 1984. The weekend effect on the distribution of stock prices: implications for option pricing. Journal of Financial Economics 13, 547–559.
Garman, M.B., Kohlhagen, S., 1983. Foreign currency option values. Journal of International Money and Finance 2, 231–237.
Hansen, P.R., Lunde, A., 2006. Realized variance and market microstructure noise (with discussion). Journal of Business and Economic Statistics 24, 127–161.
Huang, X., Tauchen, G., 2005. The relative contribution of jumps to total price variance. Journal of Financial Econometrics 3, 456–499.
Jiang, G.J., Tian, Y.S., 2005. The model-free implied volatility and its information content. Review of Financial Studies 18, 1305–1342.
Jorion, P., 1995. Predicting volatility in the foreign exchange market. Journal of Finance 50, 507–528.
Lamoureux, C.G., Lastrapes, W.D., 1993. Forecasting stock-return variance: toward an understanding of stochastic implied volatilities. Review of Financial Studies 6, 293–326.
Merton, R.C., 1973. Theory of rational option pricing. Bell Journal of Economics and Management Science 4, 141–183.
Mincer, J., Zarnowitz, V., 1969. The evaluation of economic forecasts, in J. Mincer, ed., Economic Forecasts and Expectations. NBER, New York, pp. 3–46.
Pong, S., Shackleton, M.B., Taylor, S.J., Xu, X., 2004. Forecasting currency volatility: a comparison of implied volatilities and AR(FI)MA models. Journal of Banking and Finance 28, 2541–2563.
Xu, X., Taylor, S.J., 1995. Conditional volatility and the informational efficiency of the PHLX currency options market. Journal of Banking and Finance 19, 803–821.