ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS

WORKING PAPER SERIES

The role of jumps in volatility spillovers in foreign exchange markets: meteor showers and heat waves revisited

Authors	Jérôme Lahaye, and Christopher J. Neely
Working Paper Number	2014-034D
Revision Date	April 2018
Citable Link	https://doi.org/10.20955/wp.2014.034
Suggested Citation	Lahaye, J., Neely, C.J., 2018; The role of jumps in volatility spillovers in foreign exchange markets: meteor showers and heat waves revisited, Federal Reserve Bank of St. Louis Working Paper 2014-034. URL https://doi.org/10.20955/wp.2014.034

Published In	Journal of Business & Economic Statistics
Publisher Link	https://doi.org/10.1080/07350015.2018.1512865

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

The Role of Jumps in Volatility Spillovers in Foreign Exchange Markets: Meteor Shower and Heat Waves Revisited*

Jérôme Lahaye[†], Christopher Neely[‡]

September 2016. Revised April 2018

Abstract

This paper extends the literature on geographic (heat waves) and intertemporal (meteor showers) foreign exchange volatility transmission to characterize the role of jumps and cross rate propagation. We employ multivariate heterogeneous autoregressive (HAR) models to capture the quasi-long memory properties of volatility and both Shapley-Owen R^2s and portfolio optimization exercises to quantify the contributions of information sets. We conclude that meteor showers (MS) are substantially more influential than heat waves (HW), that jumps play a modest but significant role in volatility transmission, that cross-market propagation of volatility is important, and that allowing for differential HW and MS effects and differential parameters across intraday market segments is valuable. Finally, we illustrate what types of news weaken or strengthen heat wave, meteor shower, continuous and jump patterns with sensitivity analysis.

Keywords: realized volatility, jumps, transmission, periodicity, intraday, meteor shower, heat wave, exchange rate, euro, yen, dollar

JEL Codes: C13, C14, C32, C58, F31, F37, F65, G15

*The authors would like to thank Farooq Malik, Michael McCracken, Viviana Fernandez, Carsten Trenkler, Ralf Brüggemann, Carsten Jentsch, and participants in seminars and workshops at the following institutions for helpful discussions and suggestions: the Federal Reserve Bank of St. Louis, the University of Melbourne, the University of Memphis, Midwest Econometrics Group meetings, Midwest Finance Association meetings, Victoria University, West Virginia University, the University of Louvain-la-Neuve, the University of Alabama (SNDE conference), George Washington University (Oxmetrics conference), and Sciences Po Aix-en-Provence. We are also grateful to 2 anonymous referees, the associate editor and the editor for very useful comments and suggestions. The authors retain responsibility for errors.

[†]Fordham University; jlahaye@fordham.edu; corresponding author

[‡]Assistant Vice President, Research Department, Federal Reserve Bank of St. Louis; neely@stls.frb.org

1 Introduction

Characterizing the determinants of asset price volatility and how information propagates across markets is a central goal for financial economics. Both practical forecasting applications and theoretical financial models rely on empirical volatility models. Theory requires facts about volatility to know what patterns to match in models of information release and trading patterns over time. Empirical volatility models also inform practical financial applications, such as construction of forecast confidence intervals, density forecasting, risk management, asset allocation and option pricing with time-varying volatility.

Asset price volatility tends to cluster intertemporally and to exhibit circadian patterns, including circadian autocorrelation. A number of studies have investigated whether "heat waves" (HW) or "meteor showers" (MS) dominate in exchange rate volatility transmission. Heat waves refer to the idea that geographic (i.e., circadian) autocorrelation determines volatility. A heat wave might raise volatility in New York trading on Monday and Tuesday but not in London on Tuesday morning. In contrast, meteor showers refer to the tendency of volatility to be directly temporally autocorrelated, with volatility spilling over from Asian to European to North American markets in the same day, for example.

Engle, Ito, and Lin (1990) and Ito, Engle, and Lin (1992) introduced the concepts of meteor showers and heat waves, studying the issue with generalized autoregressive heteroskedastic (GARCH) models. Baillie and Bollerslev (1991) and Hogan and Melvin (1994) extended this early research. Melvin and Peiers Melvin (2003) reinvestigated the question with a vector autoregression (VAR) for realized volatilities. Whereas Engle, Ito, and Lin (1990) found that meteor showers predominated, Melvin and Peiers Melvin (2003) argued that heat waves were more important. Thus, different approaches to volatility measurement and modeling led to different conclusions. The GARCH approach (using one observation per trading segment) pointed to the importance of meteor showers, whereas the realized volatility / VAR approach pointed to heat waves. More recently, Cai, Howorka, and Wongswan (2008) confirm Melvin and Peiers Melvin (2003) result that heat waves tend to dominate, using firm quotes from the Electronic Broking Services (EBS) trading platform, rather than indicative Reuters quotes used by Melvin and Peiers Melvin (2003).[1]

Volatility might cluster intertemporally for at least two reasons: clustering in the pattern of news and/or prolonged release of private information as heterogeneous agents trade in response to news (Engle, Ito, and Lin 1990). Clustering in news seems to be an empirical fact; many periods produce an unusual amount of news about the state of the economy.[2] Heterogeneity among agents can take the form of asymmetric information, such as private information about asset portfolios or different models of the economy.

There is empirical evidence that private information influences volatility and that news has prolonged effects on volatility. Ito, Lyons, and Melvin (1998) exploit the natural experiment of exogenous variation in Tokyo trading hours to link volatility to the release of private information while Andersen, Bollerslev, Diebold, and Vega (2003) find empirical support for a prolonged volatility response to news. Melvin and Peiers Melvin (2003) relate the theoretical microstructure literature to volatility findings particularly well.

Heat waves and meteor showers can be naturally linked to types of news. Intuitively, heat waves are more likely to occur if most or all important news that affects volatility occurs during a particular country's business day, such as the news during a week of a US debt ceiling crisis (Ito, Engle, and Lin 1992). Similarly, meteor showers will tend to predominate if autocorrelated international news is more important.

[1] Indicative Reuter quotes have the advantages of widespread availability and good informational content (Phylaktis and Chen 2009).

[2] Early researchers sometimes cited market inefficiency, in the form of bandwagon effects that might be related to technical trading, as a potential reason for volatility clustering (Ito, Engle, and Lin 1992).

Both explanations seem to be important, even for the same news events. For example, the financial crisis of 2007-2009 featured many periods of unusually important news, with policymakers and market participants reacting to each other and to events.

Our investigation incorporates methodological advances that substantially improve our understanding of volatility propagation. We study the dynamics of a covariance matrix estimated using high frequency data from four exchange rate series – the euro (EUR), Japanese yen (JPY), British pound (GBP) and Australian dollar (AUD) vs. the US dollar (USD). The use of high-frequency data substantially improves volatility modeling by allowing the calculation of "realized volatility," the sum of squared high frequency returns to estimate volatility at every instant. In theory, this statistic allows one to estimate volatility and jumps arbitrarily well, at least in the absence of microstructure noise, and therefore allows researchers to treat them as observed, rather than as noisy estimates of latent quantities (Andersen and Bollerslev 1998). We combine the Lee and Mykland (2008) and Andersen, Bollerslev, and Dobrev (2007) procedures with the threshold bipower variation of Corsi, Pirino, and Reno (2010) to detect jumps at 5-minute frequencies and evaluate their role in volatility propagation. These well-defined estimates of volatility and jumps allow more precise estimation of explanatory regressions for volatility/jumps and minimizes attenuation bias when the volatility/jump variable is a regressor. To capture the quasi-long-memory dynamics of foreign exchange volatility, we estimate heterogeneous autoregressive (HAR) models with both unconstrained linear and matrix log (i.e. *logm*) (Bauer and Vorkink 2011). We evaluate the importance of heat wave/meteor shower, continuous/jump and cross-rate propagation with both a statistical metric, the Shapley-Owen R^2, and an economic metric, the certainty equivalent return (CER) of portfolio variance minimization with varying information sets. Finally, we connect the time series model with economic ideas by using sensitivity analysis to illustrate how different types of news influence the strength of meteor shower / heat wave mechanisms.

To presage the results, these novel methods provide new insights into the issue of meteor showers and heat waves in volatility transmission, including cross-rate transmission. Specifically, the Shapley-Owen R^2 implies that meteor showers contribute 65-75% of explanatory power for realized variance versus 25-35% for heat waves. This challenges the conclusions in Cai, Howorka, and Wongswan (2008) and Melvin and Peiers Melvin (2003) that heat wave effects dominate. From a theoretical point of view, the substantial shift of the the evidence away from heat waves and toward meteor showers suggests a much greater role for international news and prolonged rebalancing after news and a lessor role for national news. Jumps have modest but positive explanatory power, accounting for 15-20% of total predictive power versus the 80-85% due to continuous variables. The relative importance of meteor showers/heat waves and continuous/jump regressors is usually similar across own and cross-series propagation. We also find substantial cross-rate jump spillovers, in contrast to the conclusions of Soucek and Todorova (2014) who found no role for jumps in spillovers between foreign exchange, the S&P 500 and commodity markets. Lack of power in the early jump estimation methods employed in Soucek and Todorova (2014) might explain this difference, as argued by Corsi, Pirino, and Reno (2010). Perhaps surprisingly, the Shapley-Owen R^2 implies that cross rate propagation — not own-rate information — accounts for the majority (60-65%) of explanatory power. This is likely largely due to the fact that, for any given exchange rate, there is information in the system from 3 other exchange rates and 6 covariance series.

An economic metric, portfolio variance minimization, implies broadly similar conclusions about the relative importance of information sets: HW/MS, continuous/jumps and own/other series. All predictors are jointly very valuable, having a daily CER of 6.64%, net of transactions costs. Meteor showers and continuously valued variables are the most valuable information sets, producing about 5.92% and 6.35% net returns, while heat waves and jump produce net CERs of 3.76% and 1.57% per annum, respectively. The portfolio optimization metric contrasts with the statistical metric in that it implies that own-series predictability produces a greater net CER than other series predictability. The net CERs for own and other information sets are 3.16% and 2.19%, respectively.

We also conclude that it is useful to account for the HW/MS and 5 intraday market segments in the portfolio variance minimization analysis. The portfolio minimization exercise shows that allowing for the HW/MS structure in the HAR produced modest CER benefits of of 0.52% per annum over a constrained specification. Also allowing coefficients to vary across segments increased those benefits to about 2.02% per annum. Although we omit full results for brevity, these benefits of allowing for HW/MS structures in the HAR terms and not constraining coefficients to be equal across equations are even substantially greater during the volatile period of the financial crisis, which was excluded from the sample. That is, the HW/MS distinction is not simply an unusual feature of the data for theorists to explain, it is a real feature with practical empirical implications.

Sensitivity analysis shows that news has the greatest and most consistent effects on patterns during the period in which the news is released. For example, US macro news has its greatest effects during the US morning and US monetary policy announcements have their greatest effects during the afternoon. The effect of news releases on statistical patterns varies by intraday segment and pattern. Particularly large shocks tend to weaken the fit of both the heat wave and meteor shower relations while smaller and more frequent shocks tend to strengthen the fit of the model.

The remainder of the paper proceeds as follows: Section 2 describes the basic covariance matrix and jump estimators used in this study, Section 3 presents data and the models for the realized covariance matrix. Section 4 presents empirical findings: results on volatility transmission and cross-market spillovers, as well as a study of the causes of meteor shower and heat waves. Section 5 concludes.

2 Covariance matrix and jump estimation

Erdemlioglu, Laurent, and Neely (2015) provide evidence supporting a Brownian semimartingale process with jumps for foreign exchange data. We use this framework and the associated realized estimators for jumps and volatility. We use a 5-minute sampling frequency because Liu, Patton, and Sheppard (2015) show that it provides a good tradeoff between accurately measuring volatility and minimizing microstructure noise. Following Lahaye, Laurent, and Neely (2011), we identify jumps with the Lee and Mykland (2008)/Andersen, Bollerslev, and Dobrev (2007) method. Dumitru and Urga (2012) show that this test is the best among a broad set of alternatives.

2.1 Jumps estimation

Define high-frequency 5-minute log-exchange-rate returns in basis points, for $n = 288$ intervals over a 24-hour trading day t, and T sample days, as follows:

$$\Delta_{j,t}P = 10000 \times (P_{j,t} - P_{j-1,t}), \quad j = 1...n, \quad t = 1, ...T, \tag{2.1}$$

where P_j is the last log-price of the interval j, and $P_{0,t} = P_{n,t-1}$.

We must remove deterministic intraday volatility patterns in returns in order to accurately identify intraday jumps. If the local volatility estimator does not account for the circadian pattern, the jump test will identify too many (few) jumps during periods of high (low) intraday periodic volatility (Boudt, Croux, and Laurent 2011b). Therefore, we adjust returns for intraday periodicity before estimating integrated volatility or jumps, and we denote these periodicity-corrected returns as $\widetilde{\Delta_{j,t}P}$. Appendix A-1.1 details our approach to correcting for periodicity, which follows Boudt, Croux, and Laurent (2011b) and Lahaye, Laurent, and Neely (2011).

The jump test statistic (Lee and Mykland 2008, Andersen, Bollerslev, and Dobrev 2007) is defined as:

$$\widetilde{J_{j,t}} = \frac{|\widetilde{\Delta_{j,t}P}|}{\sqrt{\hat{IV}_{j,t}}}, \tag{2.2}$$

where $\sqrt{\hat{IV}_{j,t}}$ is the estimated integrated volatility during period j of day t. The test statistic is half-normally distributed under the hypothesis of a Brownian semimartingale and no jumps. To reduce false jump detection, Lee and Mykland (2008) propose comparing the test statistic to critical values from the distribution of the maximum value of the statistic. We consider the maximum over the course of a day and choose a size $\alpha = 0.001$. The test detects a jump when $\widetilde{J_{j,t}} > G^{-1}(1-\alpha)S_n + C_n$, where $G^{-1}(1-\alpha)$ is the $1-\alpha$ quantile function of the standard Gumbel distribution, $S_n = \frac{1}{(2\log n)^{0.5}}$ and $C_n = (2\log n)^{0.5} - \frac{\log(\pi)+\log(\log n)}{2(2\log n)^{0.5}}$. With these parameters, we would expect to spuriously identify one jump in a 1000-day sample.

We use Corsi, Pirino, and Reno's (2010) corrected threshold bipower variation (CTBPV) to compute $\sqrt{\hat{IV}_{j,t}}$, rather than Lee and Mykland's (2008) bipower variation approach because threshold bipower variation is robust to successive jumps. Appendix A-1.2 details these power variation estimators. Unreported results actually show that the core of our results are insensitive to that choice.[3] We denote the estimated jump time series as $\hat{J}_{j,t} = \widetilde{\Delta_{j,t}P}$ when $\widetilde{J_{j,t}}$ rejects the null and $\hat{J}_{j,t} = 0$ otherwise.

2.2 Segment decomposition

To study the dynamic patterns in intraday volatility and jumps, we follow the literature in considering 5 trading segments in the following order: Asia (AS), Asia-Europe overlap (AE), Europe (EU), Europe-US overlap (ES) and the US (US) segment. We define segments as in Cai, Howorka, and Wongswan (2008), who rely on volume data, in addition to price. Their hours differ slightly from those of Melvin and Peiers Melvin (2003). For each intraday segment in our sample, we define a realized covariance matrix and estimate jumps over the segment.

Define a time index s to designate a segment (in the set $\{AS, AE, EU, ES, US\}$) on a given day. That is, s varies from 1 to $5 \times T$ (the number of segments $\times$ the number of sample days). Because there are 5 segments in each day, a given intraday segment will occur every 5th segment. For example, $s = 1$ designates the Asian (AS) segment on day 1, $s = 5$ designates the US segment on day 1, $s = 6$ the AS segment on day 2, $s = 5 \times T$ the US segment on the last sample day. Furthermore, denote n_s the number of observations in segment s. We will later use uppercase S to index the set of intraday segments.

The realized variance-covariance matrix is defined for $k = 4$ assets on a given segment, s, as follows:

$$RVC_s = R_s^{'} \times R_s, \tag{2.3}$$

[3] We use the Lee and Mykland (2008) / Andersen, Bollerslev, and Dobrev (2007) test, that identifies when intraday jumps occur. On the other hand, Corsi, Pirino, and Reno (2010) propose an aggregated jump test (i.e. the difference between realized volatility CTBPV) that do not provide jump time information as such. Therefore, our results are not directly comparable to those of Corsi, Pirino, and Reno (2010).

where R_s is an $n_s \times k$ matrix of intraday returns (computed as in Eq. 2.1) that belong to segment s.

In the next section, we describe models for the $\frac{(k \times k)+k}{2}$ (i.e. 10 with $k = 4$ assets) unique elements of RVC_s, including the 4 realized variances that occupy the diagonal of RVC_s:

$$RV_s^l = \sum_{j=1}^{n_s} \Delta_j P^{l2}, \tag{2.4}$$

where $l = 1, ..., k$ designates the asset (i.e. the exchange rate). Note that the sum on the right-hand-side is over all 5-minute elements, j, of the segment, s. The off-diagonal realized covariance between exchange rates l and m can similarly be expressed as follows:

$$RC_s^{l,m} = \sum_{j=1}^{n_s} \Delta_j P^l \times \Delta_j P^m, \tag{2.5}$$

for assets (exchange rates) $m = 1, ..., k$ and $m \neq l$.

We identify 5-minute jumps as described in Section 2.1. For segment s and exchange rate l, we aggregate jumps as follows:

$$\hat{J}_s^l = \sum_{j=1}^{n_s} \hat{J}_j^{l2}. \tag{2.6}$$

That is, a segment's estimated jump contribution is the sum of squared jump during that segment.

Finally, to test the robustness of our results to the construction of right-hand-side variables, we consider robust-to-jump alternatives to the realized variance-covariance matrix RVC_s, denoted $I\hat{V}C_s$ (estimated integrated variance-covariance matrix). We denote the diagonal elements of $I\hat{V}C_s$ as $I\hat{V}_s^l$ (estimated integrated variance for market l) and off-diagonal elements as $I\hat{C}_s^{l,m}$ (estimated integrated covariance for markets l and m). We have considered different approaches to compute $I\hat{V}C_s$, including the following: 1) setting returns identified as jumps to 0 in the computation of RVC_s[4], 2) using Boudt, Croux, and Laurent's (2011a) realized outlyingness weighted quadratic covariation, and 3) using realized bipower covariation (Barndorff-Nielsen and Shephard 2004). Our important results are not are sensitive to this choice. Therefore, we present only results for approach 1.

The next section presents data and the HAR model.

3 HAR with meteor shower/heat waves

3.1 Data and descriptive statistics

We construct covariance matrices and estimate jumps, as described in Section 2, on the EUR/USD, JPY/USD, GBP/USD and AUD/USD, using 5-minute returns, constructed from higher frequency, filtered bid data provided by Tickdata. Tickdata obtains bid quotes from three exchange archives partners using post-market batch downloads, then checks the data for errors and omissions, flags problems and suggests corrected prices.[5] We clean the exchange rate data in standard ways to remove holidays, weekends and other days with too many missing values. Lahaye, Laurent, and Neely (2011) describe the procedures. The four cleaned series cover 1803 trading days, from April 29, 2009 to July 11, 2016.

We use the 5-minute data from Tickdata to construct measures of realized (co)variance and jumps.[6]

[4] We thank a referee for suggesting this approach

[5] Additional Tickdata documentation is at https://www.tickdata.com/forex-faq/.

[6] Virtually all of our results were robust to including or excluding the year prior to the beginning of our sample, which we exclude from our main results because of its extreme volatility. Our results were also robust to not normalizing by the number of observations.

We aggregate the measures across the segments, as described in Section 2.2, and then scale them for comparability by dividing by the number of 5-minute observations in the segment. Table 1 presents descriptive statistics in basis points per 5 minutes for the elements of RVC (Eq. 2.3) and estimated jumps (Eq. 2.6) in each segment. To obtain an annualized standard deviation in percentage terms for the mean variance statistics, one would annualize by multiplying by 288*250, take the square root and divide by 100. For example, the mean EUR RV for Asia of 7.4 would convert to an annualized standard deviation of $\sqrt{(7.4 * 288 * 250)}/100 = 7.3\%$. The varying means across segments are in line with the literature. The top-left panel shows that average RV is highest during the Europe-US overlap (ES) market for all exchange rates, except the GBP. The volatility of RV (right top) is also highest during the ES for the EUR and AUD. The $I\hat{V}C$ elements (estimated integrated variances and covariances) exhibit summary statistics that are similar to those of RVC so we omit them for brevity. Not surprisingly, covariances are volatile during segments when variances are volatile. Mean squared jumps and the standard deviation of squared jumps are highest during Asia and the Europe-US overlap when averaging over all exchange rates (panel 3). The mean and standard deviation of estimated jumps are highest during Asia and the Europe-US overlap when averaging over all exchange rates (panel 3).

3.2 The HAR model

This section introduces the HAR structure that we construct to evaluate how integrated variance and jumps forecast realized volatility through HW and MS. The slowly decaying autocorrelation exhibited by foreign exchange realized volatility has prompted researchers to consider models, such as autoregressive, fractionally integrated, moving average (ARFIMA) and HAR models (see Corsi (2009)), that can fit this prominent feature of the data. We use a multivariate HAR model to capture MS and HW dynamics because it is tractable, flexible and can parsimoniously reproduce the slowly decaying autocorrelation in foreign exchange realized volatility.

Recent research has also established the importance of jumps (discontinuities) in asset prices, including their impact on volatility (see e.g. Soucek and Todorova (2014)). Because jumps are important features of volatility with different characteristics than the continuous components of volatility, it is potentially important to integrate them into a volatility spillover model in order to fully characterize the process (see Corsi, Pirino, and Reno (2010)). Therefore, we follow Andersen, Bollerslev, and Diebold (2007) in including estimated jump components in the HAR model.

In constructing the HAR model, we distinguish between HW and MS regressors. HW regressors are constructed from data observed during the same intraday period on a previous day. For example, all HW variables during the Asian segment are made up of data from previous Asian segments. MS regressors are constructed from data collected during any other period. Because there are 5 segments, all the HW regressors are constructed from variables with lags evenly divisible by 5. MS regressors are constructed from data whose lags are not evenly divisible by 5.

The regressor structure of the HAR model is complex so we will take pains to explain it clearly. We construct regressors in the HAR model from 14 series: 4 estimated IV series ($\hat{IV}_s^l$), 6 estimated integrated covariance ($\hat{IC}_s^{l,m}$) series and 4 estimated jump ($\hat{J}_s^l$) series. We construct 9 regressors from each of the 14 series. Table 2 describes the structure of the 126 ($= 9 * 14$) variable HAR regressors. Each row of Table 2 shows the 9 regressors constructed from one series while each column of Table 2 shows the same type of regressor constructed from each of the 14 series.

Consider the 9 regressors constructed from a single series. The regressors constructed from EUR estimated IV, for example, are in row 1 of Table 2. The first five regressors from any given series are the first five lags of that independent variable. Of these five lags, lags 1 through 4 are MS regressors while lag 5 is a HW regressor.

In addition, we construct four HAR variables from each series, one MS and one HW variable for weekly data and one MS and one HW variable for monthly data. The weekly (monthly) HW HAR variables are the averages of the observations in the same intraday segment over the past week (month). Columns 6 through 9 of Table 2 describe their construction. The intraday HW segments over the past week will have lags of 5, 10, 15, 20 and 25 for the past week and that set of indices is denoted by $HW(1-25)$. The HW segments for the monthly HW variable will be lags $\{5, 10, 15, ..., 95, 100\}$ and that set of indices is denoted by $HW(1-100)$. Likewise, the MS intraday segments over the past week will consist of all lags from 1 to 25 that are not HW segments. That is, the MS weekly segments will be lags $\{1, 2, 3, 4, 6, 7, ..., 24\}$ and that set will be denoted as $MS(1-25)$. The MS intraday segments over the past month will consist of all lags from 1 to 100 that are not HW segments. That is, the MS monthly segments will be lags $\{1, 2, 3, 4, 6, ..., 98, 99\}$ and that set will be denoted as $MS(1-100)$.

In each of the 5 segments of the day, we model each of the 10 (co)variance series with 9 regressors from each of the 10 covariance series and 4 jump series for a total of $9 * 14 = 126$ variable regressors plus a constant. The set of 127 explanatory variables is exactly the same within each intraday period but each intraday period will have a similarly constructed – but numerically different – set of regressors.

3.3 Estimation methods

We experimented with a wide variety of linear and nonlinear models to assess their abilities to accurately model covariance matrices at low computational cost.[7] After considering the advantages and disadvantages of each model, we choose to investigate our questions with three dissimilar models: an unconstrained linear model, a *logm* model similar to that used by Bauer and Vorkink (2011), and log-variance-fractional-logit model.[8]

The linear model is computationally convenient, provides good statistical forecasts with desirable properties but fails to enforce positive semi-definite forecast covariance matrices. The *logm* transformation enforces a positive semi-definite covariance structure but the transformed, nonlinear forecasts of the RV matrix are generally biased and often volatile. The bias can be corrected but the nonlinearly transformed variance matrix forecasts do not necessarily retain the usual in-sample properties of predictors. For example, a larger model may forecast more poorly than a smaller model if the forecasts must be nonlinearly transformed.

The unconstrained linear model regresses the 10 unique elements of the RVC matrix on the 126 HAR regressors and a constant. Expressing the regressors in Table 2 as $X_{p,t}, \quad p = 1, ..., 126$, we can write the

[7] The models included linear models, constrained and unconstrained, linear estimation of *logm* and Choleski decompositions of the volatility matrix, and direct nonlinear prediction of the volatility matrix through the inverses of these transformations.

[8] Unlike Bauer and Vorkink (2011), we do not use factors to reduce the dimensionality of the *logm* system, as we are interested in evaluating the relative contributions of different types of regressors. To do this with factors require estimating a separate group of factors for each set of regressors, which would lose much of the benefit of reduction in dimensionality and simplicity.

regression of element ij of the realized covariance matrix as follows:

$$RVC_s^{ij} = \sum_{p=1}^{126} b_p^{ij,S} X_{p,s} + b_0^{ij,S} + e_s^{ij}, \quad s = 1, 6, 11, ..., 8911^9, \tag{3.1}$$

for all i and $j \leq 4$, where S denotes the segment of the day, $S \in \{AS, AE, EU, ES, US\}$. Note that there is a different equation, with different coefficients and different regressor sets, for each segment of the day. Within each segment, each equation has the same regressors. So, there are five versions of equation 3.1, one for each segment, and the time subscripts, s, will shift forward by one period for each. That is, because we are counting in intraday segment time (not days), there is a mapping between the segment index (S) and the time/segment index (s) in the regression. The AS segment ($S = 1$) will have times, $s = 1, 6, 11, ..., 8911$, the AE segment ($S = 2$) will have times, $s = 2, 7, 12, ..., 8912$, the EU segment ($S = 3$) will have times, $s = 3, 8, 13, ..., 8913$, etc. Each coefficient in equation 3.1 has a subscript p that varies from 1 to the number of regressors (126), a superscript ij that denotes the dependent variable, and a superscript S that denotes the intraday segment with which it is associated.[10]

We construct the *logm* model forecasts in two stages. The first stage is a similar linear regression of elements of the $L_s = logm(RVC_s)$ matrix on the HAR regressors:

$$L_s^{ij} = \sum_{p=1}^{126} b_p^{ij,S} X_{p,s} + b_0^{ij,S} + e_s^{ij}, \quad s = 1, 6, 11, ..., 8911, \tag{3.2}$$

for all i and j, where L_s^{ij} is observation s of element (i, j) of the *logm* of the realized variance-covariance matrix.

In the second stage, we recovered the forecasts for RVC by taking the matrix exponential of $\hat{L}$ and — following the method of Bauer and Vorkink (2011) — pre and post multiplying $expm(\hat{L})$ by a diagonal matrix with the elements, $\sqrt{\frac{E(RV^i)}{E(R\hat{V}^i)}}$, for $i = 1, , 4$, with $E(RV^i)$ the ex-post mean realized variance, and $E(R\hat{V}^i)$ the (i, i) element of the *expm* transformation of the forecast of the *logm* matrix, i.e., it is $expm(\hat{L})$. Pre- and post-multiplying by this diagonal matrix rescales the forecast matrix to eliminate bias in variance forecasts. The resulting statistics are unbiased forecasts of the variances by construction and only very modestly biased forecasts of the covariance terms. The forecasts of RVC from the *logm* system are positive definite by construction.

The log-variance-fractional-logit model estimates the variances with a log transformation and constrains the covariances with a fractional logit specification to keep the conditional forecast correlations less than one in absolute value. Although this constraint does not strictly guarantee a positive semi-definite forecast conditional covariance matrix, it is very effective in doing so in practice. The results with this $log(variance)$-logit model are very similar to those with the logm model, so we omit the full results for brevity. It is, however, particularly useful for investigating cross market propagation and we will present those results.

4 Volatility Forecastability

4.1 How to measure forecastability?

We evaluate the importance of HW and MS effects in models that include jumps and cross-rate propagation. The existing literature has displayed the dynamic impacts of shocks with impulse response functions, but

[9]Our sample included 1803 daily observations but we lost 20 days in constructing lagged regressors. That left 1783 daily observations or 8915 segments. The last AS observation is at segment 8911.

[10]Appendix A-3 shows estimates of linear HAR coefficients for all periods.

these fail to really quantify relative contributions to predictability. There are potentially several ways to evaluate the statistical importance of explanatory power, but each has disadvantages. First, presenting coefficient estimates and their statistical significance does not lend itself easily to evaluating the effects of groups of coefficients, and the individual coefficient estimates will be marginal effects conditional on the effects of correlated regressors, not unconditional effects. Second, computing partial and/or semi-partial R^2s for groups of regressors does not effectively apportion the total explanatory power among groups of correlated regressors.

We choose two complementary strategies to evaluate the relative importance of HW versus MS, continuous variables versus jumps and cross market effects. First, to quantify the relative statistical importance of these variables, we use the Shapley-Owen R^2 measure[11] (SOR^2), sometimes called the LMG measure (Lindeman, Merenda, and Gold 1980) after its first appearance in econometrics.[12] Second, we examine the economic importance of HW/MS and continuous/jump variables with a portfolio allocation exercise for the RVC forecasts from the *logm* system.[13]

The Shapley-Owen measure of the R^2 of a group of regressors is the average improvement in R^2 for each regressor (or coalition) over all possible permutations of regressors or coalitions of regressors. The *SO* (or LMG) R^2 measure has origins in game theory and was developed to evaluate the relative importance of correlated regressors. Shapley (1953) proposed a way to apportion the gains from a cooperative game among cooperating players; Owen (1977) extended this concept to coalitions of players. Lindeman, Merenda, and Gold (1980) subsequently used the same concept to decompose goodness-of-fit among regressors and coalitions of regressors. Young (1985) shows that the Shapley Value is unique in having a set of desirable properties in a linear environment: Shapley-Owen values are efficient in that the total R^2 is distributed among the regressors or coalitions of regressors. The method treats regressors/coalitions symmetrically and, for a given refinement of predictive variables/groups, apportions contributions linearly in that a coalition's contribution is the sum of the gains of its members. For example, the impact of MS variables is the sum of the contributions of groups that use MS data. Variables with no predictive value receive a SOR^2 value of zero. But the SOR^2 becomes computationally infeasible as the number of regressor-groups (N) increases because it requires estimation of 2^N models and $N * N!$ subtractions. Therefore we estimate no more than 8 Shapley-Owen groups.

Unfortunately, the *SO* measure loses some of its desirable properties for the *logm* model's forecasts because those forecasts are nonlinear transformations (*expm*) of linear forecasts. Such a transformation means that the usual properties of forecasts might not necessarily hold. For example, forecasts from larger models might not be more highly correlated with the target variable than forecasts from smaller, nested models. Thus, the SOR^2s can be actually slightly negative, on occasion, for the bias-corrected *expm* transformations of the *logm* prediction. Despite this problem, we still report SO measures for the *logm* model because we believe that it serves as a useful check on the predictive content of types of regressors. The fact that the SOR^2s from the *logm* model are similar to and very highly correlated with those from the unconstrained linear model – where the desirable properties hold – tends to support this claim.

To estimate the contributions of HW/MS, continuous/jump regressors and cross-market predictability, we break predictive effects down in three ways: HW vs. MS, continuous (cts) vs. jump variables, own series vs. other series. As previously noted, "heat wave" variables occur in the same (lagged) period of the

[11]See Appendix A-2 for the construction of the Shapley-Owen R^2 in a simple example.
[12]See also Chevan and Sutherland (1991) and Grömping (2007).
[13]We thank the Associate Editor for this suggestion.

day as the dependent variable while "meteor shower" variables are those that occur in other periods of the data. "Jump" variables include the variable defined by exchange rate discontinuities while continuous (Cts) variables were non-jump. "Own" variables consist of series constructed exclusively from that particular exchange rate or covariance. "Other" variables came from any of the other 9 series. We classify all covariances as "other" for all variance terms (i.e., exchange rate variance) because they use information from other markets and we wished to distinguish own-market information strictly from information using other exchange rates. Thus, the definition of "Own" and "Other" depends on the dependent variable.

A comment on forecast construction is warranted here. Recall from Section 3.3 that to obtain RVC forecasts from the *logm* system, one needs to transform the predicted values of $logm(RVC)$ with the *expm* transformation. Because the *expm* transformation of the *logm* predictors interact with each other, perturbing any argument of the matrix input to *expm* will generally change all the output. To forecast some series using some subset of series in the *logm* system, one must estimate the full model using only that subset of series. That is, to forecast the EUR RV with "own" variables, we had to forecast the whole *logm* system with only EUR independent variables. We then had to repeat that process for each of the other 9 series to obtain "Own" forecasts of the whole system. This iteration was not necessary for the linear or log-variance-fractional-logit models.

4.2 Relative importance

To quantify the information content in HW and MS, we compute SOR^2s for groups of coefficients in the HAR model. The three types of effects produced a total of $2^3 = 8$ possible groupings of the 126 non-deterministic regressors: 1. $\{HW, \quad Cts, \quad Own\}$, 2. $\{HW, \quad Cts, \quad Others\}$, 3. $\{HW, \quad Jump, \quad Own\}$, 4. $\{HW, \quad Jump, \quad Others\}$, 5. $\{MS, \quad Cts, \quad Own\}$, 6. $\{MS, \quad Cts, \quad Others\}$, 7. $\{MS, \quad Jump, \quad Own\}$, 8. $\{MS, \quad Jump, \quad Others\}$. All of these groups contribute to the explanatory power of RV. We classify the exchange rate jump variables to be part of the same "own" group as the exchange rate variance. That is, the EUR jump series is part of the EUR "own" group along with the EUR variance. Covariance series, however, are not associated with "own" jumps in our model, so $\{HW, \quad Jump, \quad Own\}$ and $\{MS, \quad Jump, \quad Own\}$ provide no explanatory contribution to covariances. Therefore, covariances only have SO statistics for 6 groupings.

The two estimation methods, 5 intraday market segments, 10 dependent variables, and 6 groups of independent variables for covariances and 8 groups of independent variables for variances produce a plethora of SO results. For the sake of brevity, we initially present a limited set of Shapley-Owen statistics for the *logm* model and the Asian segment before describing broad patterns in the R^2 and SO data across models, groupings and segments of the day. To allow easier comparison of patterns in predictability across series with different R^2s, we focus on Shapley-Own ratios, that is, a group's SO R^2 as a percentage of the total R^2 for the series. These ratios sum to one by construction for each dependent variable.[14]

Table 3 shows all the 68 Shapley-Owen ratios and 10 R^2s from the *logm* (top) and unconstrained linear regression (bottom) models for the Asian period. The statistics from both models refer to predictions of the covariance series themselves. That is, the *logm* predictions have been transformed with the *expm* function and then had their bias corrected as described in Section 3.[15] The patterns are not uniform but

[14] Appendix A-3 shows Shapley-Owen ratios for all variances for all models.

[15] We compute Wald tests for the null hypotheses that the coefficients in various groups of regressors are jointly zero. Although we omit the full results for brevity, broader groups of regressors, such as HW, MS, cts and jump, were usually very significant in all or most cases. Breaking these broader groups down into smaller subgroups tended to reduce statistical

the R^2 and SO estimates from the two estimation methods are highly correlated over all segments and variables. The two predictive methods – i.e., linear and *logm* – produce R^2s that have correlations of 0.95 for covariances and 0.96 for variances while the SO ratios have correlations of 0.86 for covariances and 0.98 for variances. That is, R^2s and SO ratios that are high in the linear model also tend to be high in the *logm* model.

Table 4 provides a different perspective on the SO ratios by showing those statistics implied by the linear model for the variances for each of the 5 segments, as well as averages over all variance series and segments. We focus on the variances from the linear model in our limited space to present results. The last column (average) of Table 4 shows that the MS variables appear to provide much more explanatory power than HW; continuous variables are much more valuable than jump variables and — perhaps surprisingly — "other" variables are a bit more predictive than "own" variables. This latter result might be explained by the fact that there many more "other" independent variables than "own" independent variables for all series.[16]

Figure 1 shows how forecastability varies by method, series and intraday segment. The unconstrained linear model produces uniformly higher R^2s for all series than the transformed forecasts from the *logm* model and they both show good forecastability when one averages over all the periods of the day, as in the upper panel of Figure 1.[17] Regressions with Newey-West standard errors — omitted for brevity— which account for correlation among the R^2s from the 2 methods, indicate that these differences are statistically significant.

The upper panel of Figure 1 shows that RVs are about as forecastable, on average, as covariances. The lower panel of Figure 1 shows that the 10 series are more predictable, on average, during the middle of the global, 24-hour day, with forecastability peaking during the Asia-Europe overlap (AE) before declining over the rest of the day. Although we omit the full graphical results for brevity, all variances exhibit much greater forecastability in AE or EU, except that of the AUD/USD, which exhibits more uniform forecastability during the day.

Figure 2 characterizes mean SO ratios across segments. The upper panels compares SO ratios for each of the 8 groups, averaged over the 4 variance (left) and 6 covariance (right) series, respectively. The $\{Own, \quad jump\}$ groups always have identically zero SO ratios for covariances, as previously noted, because covariance series are not associated with jumps. The continuously valued MS groups – own and other – generally have the largest SO proportions, together making up more than half of the forecasting ability for both variance and covariance series. The two estimation methods estimate fairly similar patterns in SO ratios. The biggest difference between the unconstrained linear estimates and those from the *logm* model is that the latter estimates a higher value for $\{HW, \quad continuous, \quad other\}$ grouping for the variance series while reducing the weight on $\{MS, \quad continuous, \quad own\}$.

The upper panels of Figure 2 illustrate average SO proportions over segments of the day, which might hide significant intraday variation in the importance of the 8 categories of explanatory variables. There

significance but all groups were significant much more often than one would expect under the null. All series were usually statistically significant in predicting themselves. The Wald tests showed greater statistical significance in the AE, EU and ES periods than in the AS and US periods. Although we omit these results for brevity, we provide full results in Appendix A-3.

[16]Recall that each variance series has 18 "own" regressors, those from its own series and those from its jump series, while each covariance series just has 9 "own" regressors from its own history.

[17]The high correlations among R^2s and SO ratios and the patterns to intraday predictability (lower panel of Figure 1) are very robust to including the initial year of the financial crisis, but the R^2s are generally 5 to 20 percentage points higher with the volatile financial crisis included.

is some variability over the course of a day but still consistent patterns. The lower panels of Figure 2 illustrate that the relative importance of the 8 categories are mostly consistent over the course of a day for the linear (lower left panel) and *logm* (lower right panel) models of the variance series. The relative importance of $\{MS, \quad Cts, \quad own\}$ and $\{MS, \quad Cts, \quad other\}$ does vary somewhat between periods. $\{MS, \quad Cts, \quad other\}$ patterns do tend to be unusually important in the Asia (AS) segment and US segments. Although we omit the results for brevity, the inclusion of the volatile May-2008-April 2009 period shows patterns with more cross-segment consistency across the 5 segments, which are similar to the US patterns in the lower panel of Figure 2.

To more directly compare the relative contributions of HW vs. MS, continuous variables vs. jumps, and own vs. other variables, one can sum up the proportions of the respective categories in the upper panels of Figure 2 and compare those sums. For example, to determine the relative contribution of HW variables, one can sum the first four categories in the panels of Figure 2. This exercise implies that MS (HW) provide about 70-75% (25-30%) of the forecastability, continuous variables provide 80-85%, and own variables explain 35-40% of variances and about 32% of covariances. The results are usually similar for both the four variances and six covariances, are reasonably consistent between estimation methods, and are robust to the inclusion/exclusion of the May 2008-April 2009 period. The unconstrained linear model generally assigns a slightly greater role (a few percentage points) to continuous and own-series predictors than does the *logm* model.

In contrast to conclusions in Melvin and Peiers Melvin (2003) and Cai, Howorka, and Wongswan (2008), but in line with the seminal work of Engle, Ito, and Lin (1990), we find that MS are collectively much more important than HW. Our linear and *logm* models estimate that MS contribute about 70-75% of explanatory power for RV, versus 25-30% for HW. Given that there are four times as many MS periods than HW periods, each HW period is still much more important than a typical MS period, however, suggesting that they should be treated heterogeneously. Continuous variables contribute much more than jumps but the latter's role is not trivial. Perhaps surprisingly, there is more cross-series information than "own" series information, although any surprise should be tempered by the realization that, for any dependent variable, there is information from many "other" series but only one "own" series. All of the series are sensitive to dollar news and so they contain a lot of common information, but each individual series is a noisy estimate of that information. The lesson is that if one wishes to forecast an exchange rate variance, one should distill information from many series.

4.3 The economic value of forecastability

The literature has traditionally characterized HW/MS contributions to forecastability strictly in terms of statistical measures, such as impulse responses.[18] The previous section extended these statistical methods by introducing the Shapley-Owen measure to the task.

This section measures the economic value of varying information sets in forecasting covariance matrices. The HW-MS literature has not previously considered this very important type of metric. We assess the

[18] We forego using impulse responses in our study for two reasons. First, previous studies used a single variance series per model at each point in time, creating a natural Wold causal chain in that the observations from earlier segments neatly precede later segments, creating a natural Choleski decomposition for a vector autoregression (VAR). In contrast, we have a full covariance matrix (10 series) plus jumps at each point in time, making it unclear how to identify structural shocks in our system. Second, we believe that impulse responses are not very effective in quantifying the degree to which various types of independent variables predict a given dependent variable.

value of HW/MS and continuous/jump groups of variables by computing the certainty equivalent return (CER) for a risk-averse investor with mean-variance preferences who allocates a fixed weight portfolio across the four exchange rates each segment of each day.

The usual mean-variance optimization problem is to choose a $k \times 1$ vector of weights, w_s, to maximize the following criterion function, where Er_s is an $k \times 1$ vector of expected returns to the k assets, RVC_s is the $k \times k$ covariance matrix of those returns, and γ describes risk aversion:

$$\max_{w_s} Er_s - \frac{\gamma}{2} w_s^{'} RVC_s w_s + \lambda(w_s^{'} \mathbf{1} - 1). \tag{4.1}$$

Because we forecast the covariance matrix but not expected returns, we drop the expected returns from the criterion function. Dropping expected returns from the problem would also be justified if all expected returns were equal, in which case no choice of weight would change the expected return, $w_s^{'} Er_s$. At the end of segment s, the investor forecasts the covariance matrix in segment $s+1$ with the specified model, using full-sample forecasts with a given information set and then optimally allocates portfolio weights to the four exchange rates to minimize the variance of the portfolio, conditional on a fixed, total portfolio weight of one. In the absence of information about expected returns, the solution to the optimization problem is as follows:

$$w_s = R\hat{V}C_s^{-1} \mathbf{1} (\mathbf{1}^{'} R\hat{V}C_s^{-1} \mathbf{1})^{-1}. \tag{4.2}$$

In other words, the $k \times 1$ vector of weights is a linear combination of the inverse of the forecast covariance matrix. This result is intuitively appealing. In the absence of covariance between returns, the investor simply places higher portfolio weight on exchange rates with low variances.

The CER can be interpreted as the risk-free rate of return that an investor would pay to avoid a given risky portfolio. It is proportional to the variance of the minimized portfolio return. By comparing CER across information sets, we can measure the marginal value of information to an investor with a given level of risk aversion. We choose risk aversion (γ) to equal 3 in our exercise and transactions costs of 2 bp per unit of portfolio weight moved from one exchange rate to another. The level of risk aversion is low compared to most estimates, equal to that used in mean-variance exercises in Campbell and Thompson (2008) and less than the level of 5 used in Neely, Rapach, Tu, and Zhou (2014). The level of transactions cost is high for major currencies for the period in question (see Neely and Weller (2013)). These choices tend to conservatively estimate the economic value of information.

$$CER_p = -\frac{1}{2} \gamma w_s^{'} RVC_s w_s - \frac{cost}{2} \sum_{i=1}^{4} |w_{i,s} - w_{i,s-1}| \tag{4.3}$$

We use portfolio weights derived from the unconditional realized covariance matrix as a naïve benchmark against which to measure the value of varying information sets.[19] That is, we present CERs, net of transactions costs, for the information sets that are in excess of the CER of the portfolio optimized for the unconditional variance. In other words, if a given information set produces a net CER of 5% per annum, it means that a representative agent with CRA of 3 would be willing to pay 5% per annum for forecasts from the information set, rather than holding a portfolio whose weights minimize the unconditional variance.

The upper panel of Table 5 displays CERs in annualized percentages for in-sample optimization of the variance of a portfolio of exchange rate holdings using the *logm* HAR model and varying the information set used by the model. From left to right, the upper panel displays net annualized returns for information

[19]We also considered 1/k portfolio weights as a benchmark, but those weights implied even greater gains in CER for the econometric models — so we omit those results for brevity.

sets that use HW information only, MS information only, jump information only, continuously valued (non-jump) information only, a model that does not distinguish between HW and MS terms in constructing the weekly and monthly HAR terms, a model that pools regression coefficients over the five segments and also does not distinguish between HW and MS terms in constructing the weekly and monthly HAR terms and a model that includes all predictors.[20] The rows of the table show CERs for each of the 5 intraday periods and a rule that would use each set of forecasts in turn over the whole day. Results for each period are annualized to be comparable. We do not break down the predictors further—e.g., HW jumps and MS jumps — because we wish to directly compare HW vs. MS and continuous variables vs. jumps and the economic value of the groups are not linearly additive.

The first inference to draw from the upper panel of Table 5 is that all the information sets provide advantages over the benchmark portfolio. The continuous (non-jump) variables and MS variables each provide most of the benefit of the whole set of predictors, returning daily net CERs of 6.35% and 5.92%, respectively, over the unconditional benchmark. HW variables, by themselves, are less valuable than either the continuous (non-jump) or MS variables, but they still have considerable value, with a daily CER of 3.76%. Jump variables have the least value, with a net CER of only 1.57%. The use of all predictors permits a very respectable CER of 6.64%.

Note that the relative importance of a given information set in Table 5 is fairly consistent across daily periods. This consistency in the rank ordering of information from independent systems suggests that the differences are real. The US and EU segments show the least value to predictability while the Europe-US overlap (ES) segment shows the most. The economic value to forecasting during the ES segment mirrors that segment's very high mean RV in Table 1.

Our model departs from the traditional HAR literature in calculating HAR weekly and monthly averages separately for HW and MS data. To estimate the economic utility of this strategy, we can restrict these terms to be simply common weekly and monthly HAR averages. The column labeled "Combined HW/MS HAR terms" shows the results of the model with HW and MS weekly and monthly HAR terms combined into overall weekly and monthly HAR averages. This restriction reduces the number of parameters per segment from 127 to 99 and − comparing the columns labeled "combine HW/MS HAR terms" and "all predictors" − it costs the trading rule a modest but consistent 0.33-1.10% per annum, depending on the segment, or 0.52% over the day. This cost illustrates the economic utility of recognizing that HW and MS can have different effects.

We also consider the role of varying parameters over the segments of the day by considering the effects of pooling parameters — except constants — over the segments of the day in addition to combining the HW and MS HAR terms into overall weekly and monthly HAR. That is, instead of 5 segments that each have 127 parameters, this model uses 98 pooled parameters plus 5 constants. This pooling restriction across segments is somewhat costly, reducing the net CER from 6.64% to 4.62% per annum, on a daily basis. This conservatively estimates the gain; assuming higher risk aversion would increase the gap. It seems to be economically worthwhile to permit separate coefficients across segments and to consider separate terms for HW and MS.

We would also like to study cross-market propagation of information, but this is difficult with the *logm* model because the *expm* transformation maps all its arguments into all its outputs. As explained

[20] These information sets are sometimes but not necessarily mutually exclusive. For example, a given explanatory variable must be either a HW or MS variable but not both. On the other hand, an explanatory variable can be both a HW variable and a jump variable or neither.

previously, to estimate the effect of using only EUR/USD variance information on that series, one would have to estimate the whole system with EUR/USD variance information and then extract the forecasts of the EUR/USD variance series from the whole set of forecasts. To forecast using only "own-series" information, for example, one would have to estimate the whole model using each "own" information set and then mix forecasts from the different estimations to get a single "own-series" forecast. The resulting covariance forecasts would no longer be positive semi-definite by construction, which is a serious problem for a portfolio variance minimization exercise.

To avoid this problem, we forecast with a model that combines a log transformation of the variances with fractional-logit models that constrain the forecast correlations to be less than one in absolute value. We estimate linear HAR models for the $log(variances)$ and the fractional-logit specifications, nonlinearly transform the linear forecasts to recover forecasts of the variances and covariances, then correct the resulting estimates for bias with the Bauer-Vorkink method described for the *logm* model. Although this model does not strictly guarantee positive semi-definite covariance forecasts, in practice almost all its forecasts do satisfy this criterion.[21] Because the estimates are equation-by-equation in this model, it is fairly easy to estimate each dependent variable using only data from that series or only from other series and still nearly always get variance matrix forecasts that are positive semi-definite.

The lower panel of Table 5 displays results from this $log(variance)$-fractional-logit model in which each of the 10 covariance series is only permitted to use lagged values from the same series (own series only), a model in which each of the 10 (co)variance series are forecast only with information from the other 9 series and a model that employs all predictors for each series.[22] Perhaps surprisingly, in view of the Shapley-Owen results, the use of "own series only" predictors produces somewhat better results than the "other" series predictors, with 3.16% and 2.19% daily returns, respectively, in comparison to the 6.24% for the all predictors.[23] Again, the lesson seems to be that using information from all these series is preferable to only using a subset.

4.4 Causes of heat waves and meteor showers

We wish to investigate the source of shocks that might increase or decrease the HW and/or MS effects. One might hypothesize that temporally clustered national news — released within a country's particular working hours on successive days — creates HW effects. In contrast, temporally clustered international news or news that releases private information in the form of delayed trading creates MS effects.

To investigate what sort of news influences forecastability, we perform a variation of regression sensitivity analysis. Sensitivity analysis (a.k.a., influence analysis) is usually performed on a regression to determine if particular observations heavily influence the results. One estimates the regression, removing one observation at a time, to see if removing any particular observations cause the coefficients or other statistics to change substantially.

In our case, we calculate sensitivity series to measure the impact of macro news on statistical patterns.

[21] In the very rare cases – less than 0.1% of observations – in which the forecast covariance matrix was not positive semi-definite, we assigned the rule's portfolio to hold unconditional weights.

[22] The "all predictors" columns from the upper and lower panels differ only because the estimation methods — *logm* vs. $log(variance)$-fractional logit — differ. Note that the patterns in the sizes of segment returns in the two columns are quite similar, however.

[23] Although the results in the upper panel of Table 5 are robust to the inclusion of the financial crisis, its exclusion is important for the "own vs. other" results in the lower panel of Table 5. The use of extremely volatile financial crisis data combined with very small models produced excessively volatile and unsuccessful "own" forecasts.

That is, we remove one observation at a time and then compute the SOR^2s – not just the proportions of the R^2s presented in Table 3 – for each of 4 groups: 1) the HW variables for the RV regressors; 2) the MS variables for the RV regressors; 3) the HW variables for the jump regressors; 4) the MS variables for the jump regressors. We refer to these time series as SOR^2 sensitivity series, $R^2_{i,t}(SO)$, where i indexes the combination of the intraday segment, regressor group and the exchange rate. We examine whether news explains their variation. Specifically, we regressed the standardized changes in the SOR^2s when day t is removed from the estimation on a time series of the absolute value of standardized news surprises.[24]

$$\tilde{R}^2_{i,t}(SO) = \frac{R^2_{i,t}(SO) - \overline{R}^2_i(SO)}{\sigma_{R^2_{i,t}}} = \sum_{j=1}^{N} \beta_j \mid \frac{A_{j,t} - E(A_{j,t})}{\sigma_{A_j}} \mid + \beta_0 + \varepsilon_{i,t}, \tag{4.4}$$

where $\mid \frac{A_{j,t} - E(A_{j,t})}{\sigma_{A_j}} \mid$ is the absolute value of the standardized surprise component of the jth announcement on day t and β_0 is the constant. We estimate a version of equation 4.4 for each group's SO ratio sensitivity series for each exchange rate variance in the ES and US segments. These combinations are indexed by i in Eq. 4.4. As is usual in the announcement literature, we use survey forecasts of the announcement as the expectation to compute the surprise component and then we divide the surprise component by its standard deviation to scale the shocks.

If a particular type of news event — e.g., FOMC announcements — increases HW effects in volatility, one would expect that removing those events would reduce the SOR^2s for the HW volatility coefficients, resulting in a negative coefficients on FOMC shocks in the regression. Similarly, positive coefficients in regression 4.4 indicate that the given shock reduced estimated HW or MS effects, presumably by producing atypical volatility patterns that weaken the estimated fit.

Because we have much better data on US news and announcements than on non-US news, the estimated regressions produce much more significant results for the US and ES segments than the other three (AS, AE, and EU). Therefore, in the interests of clarity and brevity, we report results for only US and ES equations. We expect that the results for the US and ES segments would provide inference that is representative of that from the three other periods, if we had access to long samples of high-quality announcement and expectations data from Europe and Japan.

We examine 8 major US announcements that have been previously shown to have substantial effects on asset price volatility (Neely 2011). Most of these announcements are standard, including advance real Gross Domestic Product (GDP), the implicit GDP deflator, the consumer price index (CPI), the producer price index (PPI), unemployment and nonfarm payrolls (NFP). In addition, we include measures of conventional monetary policy, MP1 and ED12, developed by Kuttner (2001), Gürkaynak, Sack, and Swanson (2005) and Hausman and Wongswan (2011) and an unconventional policy measure due to Wright (2012). Fawley and Neely (2014) detail these monetary policy measures.

Table 6 shows the t-statistics, computed with Newey-West standard errors, for the coefficients from equation 4.4 for the ES and US equations, respectively. The ES segment coincides with the US morning, which is when most US macro announcements are released, while the US segment is during the US afternoon, which is when most monetary policy announcements are made. One might therefore expect US macro announcements to have greater impact during the ES segment and monetary shocks to have more influence during the US segment. The results confirm this hypothesis.

[24] We considered other specifications, including quartic time trends, an announcement indicator instead of the absolute value of the shock and day-of-the-week indicators but these alternatives either did not change the inference or did not improve the fit of the relations.

Table 6 shows that removing days of US monetary shocks have the largest and most consistent effects on the Shapley-Owen sensitivity series and that this effect is entirely during the US segment, when almost all US monetary announcements occur. Ten of 13 statistically significant MP1 shocks, which represent shocks to the short end of the yield curve, have negative coefficients, indicating that they strengthen the predictive relations. All the statistically significant coefficients on jump variables are negative, indicating that these shocks strengthen those relations. These modest conventional monetary policy shocks presumably stimulate the release of private information in a manner consistent with the rest of the data.

ED12 shocks and Wright shocks represent shocks to one-year forward and long rates, respectively. These movements tend to come from unconventional policy announcements or forward guidance and they tend to be large. It is well known that particularly large shocks present problems for volatility models that rely on autocorrelation in volatility because the very large shocks die out much more quickly than the models imply (Neely 1999, Andersen, Bollerslev, and Diebold 2007). All the statistically significant ED12 t-statistics are positive and all but one of the significant statistics on Wright shocks are positive, meaning that they tend to disrupt the estimated statistical relations. In summary, each monetary shock has a fairly consistent effect on particular statistical patterns.[25]

Among macro news, NFP, unemployment, advanced GDP and the GDP deflator series have the only statistically significant effects. NFP, in particular, has a number of statistically significant and marginal coefficients. This is consistent with the importance of these announcements in other foreign-exchange-announcement-effect studies (Lahaye, Laurent, and Neely 2011, Neely 2015). Nine of 12 significant macro announcement coefficients occur during the US morning (ES). Six of those 9 coefficients are on NFP shocks. The large but regular NFP shocks tend to have positive coefficients on continuous variables but negative coefficients on jump variables. That is, they tend to cause jumps that predict future volatility. They strengthen jump relations but may weaken the continuous relations because they are large.

In summary, news has its most significant effects on statistical patterns in the intraday trading periods in which it is announced. Monetary policy and NFP shocks strongly influence the statistical patterns, sometimes strengthening them and sometimes reducing them.

5 Conclusion

This paper extends the study of meteor showers and heat waves in foreign exchange markets by using multivariate linear and *logm* HAR models with estimated jumps and integrated volatility to characterize the forecastability of the realized variances of four exchange rate series and their six covariances. We evaluate the importance of HW/MS shower and continuous/jump information with both a statistical metric, the Shapley-Owen R^2, and an economic metric, the net CER from a portfolio variance optimization exercise with varying information sets.

From a statistical point of view, Shapley-Owen ratios imply that jumps contribute modestly to explanatory power for realized variances. Continuously valued variables contribute 80-85% of explanatory power; jumps add about 15-20%. In contrast to the previous literature, we find that meteor showers contribute much more predictive power than heat waves, 70% vs 30% of the total, although individual heat wave periods are more informative than typical meteor shower periods. Because there are 10 variance/covariance

[25] Including the very volatile financial crisis period makes many of the monetary shock coefficients statistically significant in the ES segment. The pattern of stronger effects of macro news in ES and stronger effects of monetary shocks in US remains, however. The pattern of signs on the monetary shocks during the US period also remains the same.

and 4 jump series, we find that most predictive information for realized volatility comes from "other" series rather than information from that realized volatility's own history. The linear and *logm* systems respectively imply that 60% and 65% of predictive information comes from cross-series propagation. Overall forecasting results are similar for both the four variances and six covariances and are reasonably consistent between estimation methods.

An economic metric, i.e., portfolio variance minimization, implies results about the importance of these information sets that are broadly consistent with the inference from the statistical metric. All predictors are jointly very valuable, having a CER of 6.64% annualized over all periods, net of transactions costs. Among all predictors, the meteor showers and continuously valued variables are the most valuable information sets, producing net CERs of 5.92 and 6.35%, respectively. Heat waves are the next most valuable, producing a net CER of 3.76%, with jumps producing a more modest but still respectable net CER of 1.57%. In contrast to the statistical results, using only own-series information to forecast the covariance matrix provides somewhat more information than using only other series. The two information sets have net CERs of 3.16% and 2.19%, respectively. All of these information sets appear to be quite valuable in isolation and even more valuable together.

The MS domination of HW effects directly challenges the conclusions in Melvin and Peiers Melvin (2003) and Cai, Howorka, and Wongswan (2008). We attribute the disparity between our conclusions and theirs to the fact that these previous authors used VAR impulse reponse functions to informally judge effects.

More generally, our analysis points to the importance of accounting for the HW/MS structure in the HAR and of estimating the 5 intraday segments as different processes. That is, the portfolio minimization exercise shows that accounting for the HW/MS structure in the HAR produces modest CER benefits of 0.33 to 1.10% points per annum, but allowing coefficients to vary across segments increases those benefits to about 2.02% per annum.

Our work qualifies some previous equity market research on the importance of jumps in volatility propagation. Specifically, Andersen, Bollerslev, and Diebold (2007) and Soucek and Todorova (2014) found no role for jumps in volatility prediction, within and across markets, respectively. Our approach supports significant cross-market forex jump spillovers. As Corsi, Pirino, and Reno (2010) argue, this may be due to desirable properties of the corrected threshold bipower variation employed in our jump estimation. Our results do not fully support a view of jumps as reflecting quick and definitive information incorporation. It appears that when news causes a jump, markets may take many hours to fully resolve uncertainty. Otherwise, it may be that news causing jumps may be correlated with future surprises.

This study also used sensitivity analysis to investigate how news influences statistical patterns from heat wave/meteor showers and continuous/jump variables. All announcements have the largest and most consistent effects during the segments in which they occur. That is, macro news have a greater impact in the US morning, when US macro announcements are made, while US monetary policy news has very large and consistent influences on statistical patterns in the afternoon, when FOMC announcements are generally made. In line with their importance in other studies, NFP shocks have the most statistically significant effects on forecasting patterns among macro news releases (Lahaye, Laurent, and Neely 2011, Neely 2015). Also consistent with previous research, US monetary policy shocks have the largest and most consistent impact on volatility patterns. While shocks to the short end of the yield curve strengthen the predictive relations, presumably because they increase volatility in a consistent manner, monetary shocks to the one-year ahead forward rates and long end tend to weaken statistical relations. Past research has shown that large shocks reduce the fit of volatility models while smaller shocks improve it by reinforcing patterns created by other news (Neely 1999, Andersen, Bollerslev, and Diebold 2007).

In summary, this paper has integrated estimated jumps into multivariate linear and *logm* HAR models of cross-rate volatility propagation. Results from these models have established the importance of distinguishing between heat waves and meteor showers in volatility propagation and have overturned previous conclusions about their relative importance, showing that meteor showers are much more important than heat waves by both statistical and economic metrics. Our study also has shed light on the impact of news on volatility patterns.

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, and F. X. Diebold (2007): "Roughing It Up: Including Jump Components in the Measurement, Modelling and Forecasting of Return Volatility," *Review of Economics and Statistics*, 89, 701–720.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and C. Vega (2003): "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange," *The American Economic Review*, 93, 38–62.

Andersen, T. G., T. Bollerslev, and D. Dobrev (2007): "No-arbitrage Semi-martingale Restrictions for Continous-time Volatility Models Subject to Leverage Effects, Jumps and i.i.d. Noise: Theory and Testable Distributional Implications," *Journal of Econometrics*, 138, 125–180.

Baillie, R. T., and T. Bollerslev (1991): "Intra-Day and Inter-Market Volatility in Foreign Exchange Rates," *Review of Economic Studies*, 58, 565–585.

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

Bauer, G. H., and K. Vorkink (2011): "Forecasting Multivariate Realized Stock Market Volatility," *Journal of Econometrics*, 160, 93–101.

Boudt, K., C. Croux, and S. Laurent (2011a): "Outlyingness Weighted Quadratic Covariation," *Journal of Financial Econometrics*, 9, 657–684.

——— (2011b): "Robust Estimation of Intraweek Periodicity in Volatility and Jump Detection," *Journal of Empirical Finance*, 18, 353–367.

Cai, F., E. Howorka, and J. Wongswan (2008): "Informational Linkages Across Trading Regions: Evidence from Foreign Exchange Markets," *Journal of International Money and Finance*, 27, 1215–1243.

Campbell, J. Y., and S. B. Thompson (2008): "Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average?," *Review of Financial Studies*, 21, 1509–1531.

Chevan, A., and M. Sutherland (1991): "Hierarchical Partitioning," *The American Statistician*, 45, 90–96.

Corsi, F. (2009): "A Simple Approximate Long-Memory Model of Realized Volatility," *Journal of Financial Econometrics*, 7, 174–196, doi:10.1093/jjfinec/nbp001.

Corsi, F., D. Pirino, and R. Reno (2010): "Threshold Bipower Variation and the Impact of Jumps on Volatility Forecasting," *Journal of Econometrics*, 159, 276–288.

Dumitru, A. M., and G. Urga (2012): "Identifying Jumps in Financial Assets: A Comparison Between Nonparametric Jump Tests," *Journal of Business and Economic Statistics*, 30, 242–255.

Engle, R. F., T. Ito, and W. L. Lin (1990): "Meteor Showers or Heat Waves? Heteroskedastic Intra-daily Volatility in the Foreign Exchange Market," *Econometrica*, 58, 525–542.

Erdemlioglu, D., S. Laurent, and C. J. Neely (2015): "Which Continuous-time Model is Most Appropriate for Exchange Rates?," *Journal of Banking and Finance*, 61, S256–S268.

Fawley, B., and C. J. Neely (2014): "The Evolution of Federal Reserve Policy and the Impact of Monetary Policy Surprises on Asset Prices," *Federal Reserve Bank of St. Louis Review*, 96, 73–109.

Grömping, U. (2007): "Estimators of Relative Importance in Linear Regression Based on Variance Decomposition," *The American Statistician*, 61, 139–147.

Gürkaynak, R. S., B. P. Sack, and E. T. Swanson (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 1, 55–93.

Hausman, J., and J. Wongswan (2011): "Global Asset Prices and FOMC Announcements," *Journal of International Money and Finance*, 30, 547–71.

Hogan, K., and M. Melvin (1994): "Sources of Meteor Showers and Heat Waves in the Foreign Exchange Market," *Journal of International Economics*, 37, 239–247.

Ito, T., R. F. Engle, and W. L. Lin (1992): "Where Does the Meteor Shower Come From? The Role of Stochastic Policy Coordination," *Journal of International Economics*, 32, 221–240.

Ito, T., R. K. Lyons, and M. T. Melvin (1998): "Is There Private Information in the FX Market? The Tokyo Experiment," *Journal of Finance*, 53, 1111–1130.

Kuttner, K. N. (2001): "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market," *Journal of Monetary Economics*, 47, 523–44.

Lahaye, J., S. Laurent, and C. J. Neely (2011): "Jumps, Cojumps and Macro Announcements," *Journal of Applied Econometrics*, 26, 893–921.

Lee, S. S., and P. A. Mykland (2008): "Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics," *Review of Financial Studies*, 21, 2535–2563.

Lindeman, R., P. Merenda, and R. Gold (1980): *Introduction to Bivariate and Multivariate Analysis*. Scott Foresman, Glenview, IL.

Liu, L., A. Patton, and K. Sheppard (2015): "Does Anything Beat 5-Minute RV? A Comparison of Realized Measures Across Multiple Asset Classes," *Journal of Econometrics*, 187, 293–311.

Melvin, M., and B. Peiers Melvin (2003): "The Global Transmission of Volatility in the Foreign Exchange Market," *The Review of Economics and Statistics*, 85, 670–679.

Neely, C. J. (1999): "Target Zones and Conditional Volatility: The Role of Realignments," *Journal of Empirical Finance*, 6, 177–192.

——— (2011): "A Survey of Announcement Effects on Foreign Exchange Volatility and Jumps," *Federal Reserve Bank of St. Louis Review*, 93, 361–385.

——— (2015): "Unconventional Monetary Policy Had Large International Effects," *Journal of Banking and Finance*, 52, 101–111.

Neely, C. J., D. E. Rapach, J. Tu, and G. Zhou (2014): "Forecasting the Equity Risk Premium: The Role of Technical Indicators," *Management Science*, 60, 1772–1791.

Neely, C. J., and P. A. Weller (2013): "Lessons From the Evolution of Foreign Exchange Trading Strategies," *Journal of Banking and Finance*, 37, 3783–3798.

Owen, G. (1977): "Values of Games with a Priori Unions," in *Essays in Mathematical Economics and Game Theory*, ed. by R. Henn, and O. Moeschlin, pp. 76–88. Springer, Berlin.

Phylaktis, K., and L. Chen (2009): "Price Discovery in Foreign Exchange Markets: A Comparison of Indicative and Actual Transaction Prices," *Journal of Empirical Finance*, 16, 640–654.

Shapley, L. (1953): "A Value for n-Person Games," reprinted in *The Shapley Value: Essays in Honor of Lloyd S. Shapley*, ed. A. Roth, Cambridge: Cambridge University Press.

Soucek, M., and N. Todorova (2014): "Realized Volatility Transmission: The Role of Jumps and Leverage Effects," *Economics Letters*, 122, 111–115.

Wright, J. H. (2012): "What does Monetary Policy do to Long-term Interest Rates at the Zero Lower Bound?," *The Economic Journal*, 122, 447–66.

Young, H. P. (1985): "Monotonic Solutions of Cooperative Games," *International Journal of Game Theory*, 14 (2), 65–72.

Table 1: Descriptive statistics for the realized covariance matrix and estimated jumps series

	Mean					Standard deviation				
	AS	AE	EU	ES	US	AS	AE	EU	ES	US
EUR RV	7.4	20.1	15.6	27.8	11.3	12.1	19.2	16.4	31.9	22.8
JPY RV	14.5	16.9	12.5	24.2	11.9	52.5	23.4	13.3	32.8	26.0
GBP RV	7.5	21.5	18.1	20.8	8.2	70.6	31.0	24.3	21.2	14.8
AUD RV	20.9	26.5	20.7	35.1	18.9	25.1	24.0	20.3	39.7	27.9
mean	12.6	21.3	16.7	27.0	12.6	40.0	24.4	18.6	31.4	22.9
(EUR, JPY) RC	-0.9	-1.6	-1.4	-7.2	-2.8	12.2	9.6	5.7	18.9	11.9
(EUR, GBP) RC	4.1	10.4	7.6	15.0	6.1	16.6	15.5	9.1	15.5	15.2
(EUR, AUD) RC	5.9	12.7	9.3	18.3	8.4	10.3	14.8	12.0	23.0	17.0
(JPY, GBP) RC	0.4	0.0	-0.5	-4.1	-1.6	30.3	9.8	5.3	11.8	9.1
(JPY, AUD) RC	0.4	-0.6	-0.9	-3.2	-1.3	21.0	12.7	7.3	20.8	14.6
(GBP, AUD) RC	5.2	11.0	7.8	14.6	6.5	17.2	14.1	9.5	18.2	12.9
mean	2.5	5.3	3.7	5.5	2.5	17.9	12.8	8.1	18.0	13.5
EUR $\hat{J}$	0.3	0.4	0.4	2.8	1.6	1.9	4.8	5.6	19.1	14.8
JPY $\hat{J}$	2.8	0.9	0.3	2.7	1.8	38.6	13.4	3.2	17.3	15.6
GBP $\hat{J}$	0.8	0.6	1.5	0.9	0.8	19.2	6.8	13.2	8.4	8.7
AUD $\hat{J}$	3.0	0.2	0.6	2.0	1.8	11.9	4.1	6.2	21.6	11.8
mean	1.7	0.5	0.7	2.1	1.5	17.9	7.3	7.1	16.6	12.7

Note: The panels show statistics on segments' realized covariance matrix (RVC_s) elements (Eq. 2.3), as well as segments' estimated jumps $\hat{J}_{s,l}$ (Eq. 2.6) per 5-min period (i.e. RVC_s and $\hat{J}_{s,l}$ are divided by the number of observations in their corresponding segments n_s). Original 5-minute returns are in basis points (Eq. 2.1). The left panel shows means while the right-hand panel show standard deviations. To convert mean 5-minute variances in b.p. to annualized standard deviations in percentages, one should multiply by the product of the number of 5 minute periods per day and the number of business days in a year, divide by 10,000 and take the square root. So, a mean variance of 7.4 would imply an annualized standard deviation of 7.3% ($\sqrt{7.4 * 288 * 250/10000}$).

Table 2: Table of regressors

		1	2	3	4	5	6	7	8	9
		MS	MS	MS	MS	HW	MS	HW	MS	HW
		Lag 1	Lag 2	Lag 3	Lag 4	Lag 5	weekly	weekly	monthly	monthly
Series:										
Continuous										
$\hat{IV}^{E}_{s}$	1	$\hat{IV}^{E}_{s-1}$	$\hat{IV}^{E}_{s-2}$	$\hat{IV}^{E}_{s-3}$	$\hat{IV}^{E}_{s-4}$	$\hat{IV}^{E}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{IV}^{E}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{IV}^{E}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{IV}^{E}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{IV}^{E}_{s-i}/20$
$\hat{IC}^{E,J}_{s}$	2	$\hat{IC}^{E,J}_{s-1}$	$\hat{IC}^{E,J}_{s-2}$	$\hat{IC}^{E,J}_{s-3}$	$\hat{IC}^{E,J}_{s-4}$	$\hat{IC}^{E,J}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{IC}^{E,J}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{IC}^{E,J}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{IC}^{E,J}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{IC}^{E,J}_{s-i}/20$
$\hat{IC}^{E,G}_{s}$	3	$\hat{IC}^{E,G}_{s-1}$	$\hat{IC}^{E,G}_{s-2}$	$\hat{IC}^{E,G}_{s-3}$	$\hat{IC}^{E,G}_{s-4}$	$\hat{IC}^{E,G}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{IC}^{E,G}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{IC}^{E,G}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{IC}^{E,G}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{IC}^{E,G}_{s-i}/20$
$\hat{IC}^{E,A}_{s}$	4	$\hat{IC}^{E,A}_{s-1}$	$\hat{IC}^{E,A}_{s-2}$	$\hat{IC}^{E,A}_{s-3}$	$\hat{IC}^{E,A}_{s-4}$	$\hat{IC}^{E,A}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{IC}^{E,A}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{IC}^{E,A}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{IC}^{E,A}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{IC}^{E,A}_{s-i}/20$
$\hat{IV}^{J}_{s}$	5	$\hat{IV}^{J}_{s-1}$	$\hat{IV}^{J}_{s-2}$	$\hat{IV}^{J}_{s-3}$	$\hat{IV}^{J}_{s-4}$	$\hat{IV}^{J}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{IV}^{J}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{IV}^{J}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{IV}^{J}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{IV}^{J}_{s-i}/20$
$\hat{IC}^{J,G}_{s}$	6	$\hat{IC}^{J,G}_{s-1}$	$\hat{IC}^{J,G}_{s-2}$	$\hat{IC}^{J,G}_{s-3}$	$\hat{IC}^{J,G}_{s-4}$	$\hat{IC}^{J,G}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{IC}^{J,G}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{IC}^{J,G}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{IC}^{J,G}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{IC}^{J,G}_{s-i}/20$
$\hat{IC}^{J,A}_{s}$	7	$\hat{IC}^{J,A}_{s-1}$	$\hat{IC}^{J,A}_{s-2}$	$\hat{IC}^{J,A}_{s-3}$	$\hat{IC}^{J,A}_{s-4}$	$\hat{IC}^{J,A}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{IC}^{J,A}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{IC}^{J,A}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{IC}^{J,A}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{IC}^{J,A}_{s-i}/20$
$\hat{IV}^{G}_{s}$	8	$\hat{IV}^{G}_{s-1}$	$\hat{IV}^{G}_{s-2}$	$\hat{IV}^{G}_{s-3}$	$\hat{IV}^{G}_{s-4}$	$\hat{IV}^{G}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{IV}^{G}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{IV}^{G}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{IV}^{G}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{IV}^{G}_{s-i}/20$
$\hat{IC}^{G,A}_{s}$	9	$\hat{IC}^{G,A}_{s-1}$	$\hat{IC}^{G,A}_{s-2}$	$\hat{IC}^{G,A}_{s-3}$	$\hat{IC}^{G,A}_{s-4}$	$\hat{IC}^{G,A}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{IC}^{G,A}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{IC}^{G,A}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{IC}^{G,A}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{IC}^{G,A}_{s-i}/20$
$\hat{IV}^{A}_{s}$	10	$\hat{IV}^{A}_{s-1}$	$\hat{IV}^{A}_{s-2}$	$\hat{IV}^{A}_{s-3}$	$\hat{IV}^{A}_{s-4}$	$\hat{IV}^{A}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{IV}^{A}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{IV}^{A}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{IV}^{A}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{IV}^{A}_{s-i}/20$
Jumps										
$\hat{J}^{E}_{s}$	11	$\hat{J}^{E}_{s-1}$	$\hat{J}^{E}_{s-2}$	$\hat{J}^{E}_{s-3}$	$\hat{J}^{E}_{s-4}$	$\hat{J}^{E}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{J}^{E}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{J}^{E}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{J}^{E}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{J}^{E}_{s-i}/20$
$\hat{J}^{J}_{s}$	12	$\hat{J}^{J}_{s-1}$	$\hat{J}^{J}_{s-2}$	$\hat{J}^{J}_{s-3}$	$\hat{J}^{J}_{s-4}$	$\hat{J}^{J}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{J}^{J}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{J}^{J}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{J}^{J}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{J}^{J}_{s-i}/20$
$\hat{J}^{G}_{s}$	13	$\hat{J}^{G}_{s-1}$	$\hat{J}^{G}_{s-2}$	$\hat{J}^{G}_{s-3}$	$\hat{J}^{G}_{s-4}$	$\hat{J}^{G}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{J}^{G}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{J}^{G}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{J}^{G}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{J}^{G}_{s-i}/20$
$\hat{J}^{A}_{s}$	14	$\hat{J}^{A}_{s-1}$	$\hat{J}^{A}_{s-2}$	$\hat{J}^{A}_{s-3}$	$\hat{J}^{A}_{s-4}$	$\hat{J}^{A}_{s-5}$	$\sum_{i\in MS(1-25)} \hat{J}^{A}_{s-i}/20$	$\sum_{i\in HW(1-25)} \hat{J}^{A}_{s-i}/5$	$\sum_{i\in MS(1-100)} \hat{J}^{A}_{s-i}/80$	$\sum_{i\in HW(1-100)} \hat{J}^{A}_{s-i}/20$

Note: We construct 126 variable regressors in the HAR model from 14 series: 10 (co)variance series and 4 jump series. From each of the 14 series, we construct 9 regressors. The first five are simple lags from 1 to 5 of the independent variable. Of these, lags 1 through 4 are meteor shower regressors while lag 5 is a heat wave regressor. In addition, we construct four HAR variables, one MS and one HW variable for weekly data and one MS and one HW variable for monthly data. The weekly (monthly) HW HAR variables are the averages of the observations in the same intraday segment over the past week (month). The intraday HW segments over the past week will have lags of 5, 10, 15, 20 and 25 for the past week and that set of indices is denoted by HW(1-25). The HW segments for the monthly HW variable will be lags $5, 10, 15, ..., 95, 100$ and that set of indices is denoted by HW(1-100). The MS intraday segments over the past week will consist of all lags from 1 to 25 that are not HW segments. That is, the MS weekly segments will be lags $1, 2, 3, 4, 6, 7, ...24$ and that set will be denoted as MS(1-25). The MS intraday segments over the past month will consist of all lags from 1 to 100 that are not HW segments. That is, the MS monthly segments will be lags $1, 2, 3, 4, 6, ..., 98, 99$ and that set will be denoted as MS(1-100). The table above shows all 126 variable regressors and labels the columns as MS or HW variables. Two thirds of the explanatory variables are MS. The top panel of the table shows the 90 continuous regressors, which are constructed from the 10 (co)variance series while the lower panel shows the 36 jump regressors constructed from the 4 jump series. The EUR, JPY, GBP and AUD variances are denoted by E, J, G and A, respectively. The six covariances from those four exchange rates are similarly denoted by (E, J), (E, G), (E, A), (J, G), (J, A) and (G, A).

Table 3: Shapley-Owen ratios for the linear and *logm* model during the Asian segment

logm model	realized variance				realized covariance					
	EUR	JPY	GBP	AUD	(EUR, JPY)	(EUR, GBP)	(EUR, AUD)	(JPY, GBP)	(JPY, AUD)	(GBP, AUD)
$\{HW, Cts, Own\}$ % of total	-0.002	0.009	0.082	-0.025	0.041	0.022	0.091	0.001	0.004	0.037
$\{HW, Cts, Others\}$ % of total	0.237	0.250	0.072	0.198	0.211	0.217	0.239	0.287	0.288	0.230
$\{HW, Jump, Own\}$ % of total	0.008	0.011	-0.004	-0.002						
$\{HW, Jump, Others\}$ % of total	0.043	0.096	0.007	0.030	0.112	0.088	0.058	0.091	0.106	0.052
$\{MS, Cts, Own\}$ % of total	0.233	0.127	0.391	0.248	0.191	0.120	0.217	0.138	0.096	0.117
$\{MS, Cts, Others\}$ % of total	0.411	0.466	0.454	0.501	0.414	0.518	0.345	0.492	0.506	0.544
$\{MS, Jump, Own\}$ % of total	0.005	0.021	0.001	0.015						
$\{MS, Jump, Others\}$ % of total	0.066	0.019	-0.004	0.036	0.031	0.036	0.051	-0.009	-0.001	0.020
total R^2	0.234	0.120	0.255	0.226	0.129	0.135	0.269	0.283	0.127	0.144
unconstrained linear model	realized variance				realized covariance					
	EUR	JPY	GBP	AUD	(EUR, JPY)	(EUR, GBP)	(EUR, AUD)	(JPY, GBP)	(JPY, AUD)	(GBP, AUD)
$\{HW, Cts, Own\}$ % of total	0.061	0.029	0.024	0.090	0.025	0.025	0.100	0.016	0.032	0.033
$\{HW, Cts, Others\}$ % of total	0.130	0.094	0.082	0.141	0.123	0.093	0.144	0.099	0.106	0.100
$\{HW, Jump, Own\}$ % of total	0.004	0.006	0.026	0.003						
$\{HW, Jump, Others\}$ % of total	0.027	0.027	0.007	0.024	0.035	0.029	0.025	0.033	0.028	0.026
$\{MS, Cts, Own\}$ % of total	0.162	0.181	0.174	0.172	0.135	0.074	0.172	0.214	0.122	0.109
$\{MS, Cts, Others\}$ % of total	0.559	0.609	0.654	0.519	0.643	0.740	0.516	0.602	0.680	0.692
$\{MS, Jump, Own\}$ % of total	0.015	0.022	0.009	0.010						
$\{MS, Jump, Others\}$ % of total	0.044	0.033	0.024	0.041	0.039	0.038	0.043	0.035	0.033	0.040
total R^2	0.414	0.284	0.341	0.389	0.354	0.353	0.460	0.343	0.361	0.362

Note: The upper panel shows SO ratios and R^2s from the *logm* model while the lower panel shows analogous statistics from the unconstrained linear model. A SO ratio is the proportion of the total R^2 attributed to the specific group. The SO ratios sum to one for each series (column) and method of estimation (upper and lower panels).

Table 4: Linear model Shapley-Owen ratios for variances for all segments

		AS				AE				EU			
		EUR	JPY	GBP	AUD	EUR	JPY	GBP	AUD	EUR	JPY	GBP	AUD
$\{HW, Cts, Own\}$	% of total	0.061	0.029	0.024	0.090	0.111	0.086	0.051	0.148	0.147	0.113	0.108	0.162
$\{HW, Cts, Other\}$	% of total	0.130	0.094	0.082	0.141	0.112	0.092	0.050	0.167	0.165	0.082	0.118	0.175
$\{HW, Jump, Own\}$	% of total	0.004	0.006	0.026	0.003	0.009	0.005	0.001	0.002	0.004	0.012	0.002	0.003
$\{HW, Jump, Other\}$	% of total	0.027	0.027	0.007	0.024	0.014	0.020	0.002	0.013	0.007	0.006	0.012	0.007
$\{MS, Cts, Own\}$	% of total	0.162	0.181	0.174	0.172	0.297	0.297	0.270	0.278	0.259	0.460	0.263	0.288
$\{MS, Cts, Other\}$	% of total	0.559	0.609	0.654	0.519	0.304	0.284	0.281	0.322	0.367	0.219	0.284	0.314
$\{MS, Jump, Own\}$	% of total	0.015	0.022	0.009	0.010	0.014	0.158	0.217	0.015	0.011	0.047	0.120	0.015
$\{MS, Jump, Other\}$	% of total	0.044	0.033	0.024	0.041	0.140	0.057	0.129	0.055	0.041	0.061	0.092	0.036
R2		0.414	0.284	0.341	0.389	0.622	0.351	0.809	0.620	0.399	0.565	0.464	0.582
		ES				US				Average			
		EUR	JPY	GBP	AUD	EUR	JPY	GBP	AUD				
$\{HW, Cts, Own\}$	% of total	0.165	0.087	0.127	0.146	0.063	0.082	0.078	0.128	0.100			
$\{HW, Cts, Other\}$	% of total	0.207	0.109	0.132	0.192	0.110	0.101	0.095	0.164	0.126			
$\{HW, Jump, Own\}$	% of total	0.010	0.022	0.017	0.009	0.007	0.026	0.007	0.005	0.009			
$\{HW, Jump, Other\}$	% of total	0.053	0.026	0.007	0.041	0.013	0.017	0.015	0.028	0.018			
$\{MS, Cts, Own\}$	% of total	0.204	0.213	0.233	0.241	0.112	0.306	0.311	0.269	0.249			
$\{MS, Cts, Other\}$	% of total	0.280	0.395	0.282	0.290	0.656	0.430	0.318	0.378	0.387			
$\{MS, Jump, Own\}$	% of total	0.033	0.107	0.119	0.050	0.015	0.019	0.104	0.004	0.055			
$\{MS, Jump, Other\}$	% of total	0.049	0.041	0.083	0.031	0.023	0.019	0.072	0.024	0.055			
R^2		0.228	0.399	0.627	0.429	0.255	0.252	0.400	0.417	0.442			

Note: The five panels show SO ratios and R^2s for variance series implied by the linear model for each of the 5 daily segments. A SO ratio is the proportion of the total R^2 attributed to the specific group. The SO ratios sum to one for each series (column).

Table 5: Certainty equivalent returns from forecasts with varying information sets

	logm model						
	HW	MS	Jumps	Cts	Combine HW/MS HAR terms	Pooled, combine HW/MS HAR terms	All predictors
AS	2.40	6.56	2.32	7.26	6.99	3.35	7.68
AE	6.10	7.59	1.66	8.76	7.62	8.03	8.71
EU	4.13	5.18	0.66	5.37	5.17	4.72	5.56
ES	7.95	10.59	2.16	10.69	10.82	9.33	11.36
US	3.74	4.36	1.07	4.65	4.40	4.12	4.73
Daily	3.76	5.92	1.57	6.35	6.12	4.62	6.64

	fractional logit model		
	Own Own	Other Other	All predictors
AS	4.47	2.28	7.32
AE	2.90	3.31	8.05
EU	2.12	1.34	4.82
ES	4.35	4.34	10.58
US	2.68	1.45	4.59
Daily	3.16	2.19	6.24

Note: The upper (lower) panel of the table displays certainty equivalent returns in annualized percentages for in-sample optimization of the variance of a portfolio of exchange rate holdings using the *logm* ($log(variance)$-fractional logit) HAR model and varying the information set used by the model. The columns of the upper panel display results, from left to right, for information sets that use heat-wave information only, meteor-shower information only, jump information only, continuously valued (non-jump) information only, a model that does not distinguish between HW and MS terms in constructing the weekly and monthly HAR terms, a model that pools regression coefficients over the five segments and also does not distinguish between HW and MS terms in constructing the weekly and monthly HAR terms and a model that includes all 126 predictors. These information sets are sometimes but not necessarily mutually exclusive. For example, a given explanatory variable must be either a heat wave or meteor shower variable but not both. On the other hand, an explanatory variable can be both a heat wave variable and a jump variable. The rows of the panel display results from the 5 intraday segments and a weighted average measure for a daily return. The lower panel shows the analogous statistics from a model that combines a log transformation of the variances with fractional logit models that constrain the forecast correlations to be less than one in absolute value. From left to right, the lower panel displays results from a model in which each of the 10 covariance series is only permitted to use lagged values from the same series (own series only), a model in which each of the 10 covariance series are forecast only with information from the other 9 series and a model that employs all predictors for each series. The "all predictors" columns from the respective panels differ only because the estimation methods — *logm* vs. $log(variance)$-fractional logit — differ.

Table 6: Sensitivity analysis of factors strengthening or weakening patterns

ES segment									
	Adv GDP	GDP Deflator	CPI	PPI	Unemployment	NFP	Wright	MP1	ED12
Eur HW Cts	2.45	-1.48	1.19	-0.37	-1.18	-0.20	0.22	-0.02	0.14
JPY HW Cts	0.11	-0.15	-0.08	-0.38	-1.09	0.37	0.06	0.07	0.05
GBP HW Cts	0.21	-0.06	-0.39	-0.21	-0.04	-0.46	0.01	-0.01	-0.09
AUD HW Cts	0.00	-0.15	0.09	-0.36	-0.17	0.04	-0.02	0.01	-0.11
Eur HW J	0.65	-0.56	0.84	0.02	0.93	-2.83	0.28	-0.24	0.06
JPY HW J	-0.07	-0.18	0.30	-0.16	2.34	-2.28	0.02	-0.07	0.15
GBP HW J	0.19	-0.13	-0.14	-0.44	0.86	-5.11	0.33	0.02	-0.56
AUD HW J	0.04	-0.07	-0.21	-0.50	1.05	-6.19	0.08	-0.05	-0.19
Eur MS Cts	1.43	-0.97	-1.21	-0.29	-1.47	2.34	0.49	-0.14	0.34
JPY MS Cts	0.52	0.16	0.96	0.09	-0.44	3.67	0.09	-0.23	0.01
GBP MS Cts	0.21	0.15	-0.28	0.18	0.40	1.54	-0.07	-0.12	0.13
AUD MS Cts	0.15	0.03	-1.70	-0.10	0.35	1.68	0.04	0.15	-0.04
Eur MS J	1.16	-0.62	-0.01	-0.09	1.04	-1.06	-0.14	0.62	-0.30
JPY MS J	0.18	0.00	-0.08	-0.15	-2.46	-1.57	-0.01	-0.03	-0.01
GBP MS J	0.05	0.11	0.20	0.15	0.05	0.69	-0.02	0.04	0.01
AUD MS J	0.26	0.06	0.06	-0.07	-0.01	0.90	0.01	0.02	-0.12

US segment									
	Adv GDP	GDP Deflator	CPI	PPI	Unemployment	NFP	Wright	MP1	ED12
Eur HW Cts	-1.35	0.69	-0.17	-0.14	-0.13	-0.17	-1.27	-9.59	20.93
JPY HW Cts	0.07	-0.09	-0.06	-0.33	-0.05	-0.19	0.27	-1.35	-0.37
GBP HW Cts	-1.40	0.65	-0.40	-0.17	-0.10	-0.06	-1.13	-9.29	21.08
AUD HW Cts	-2.69	-2.18	-0.95	-0.71	-0.12	-0.39	10.22	8.17	11.16
Eur HW J	-1.46	0.52	-0.53	-0.29	-0.08	-0.45	-1.23	-10.10	21.46
JPY HW J	-0.05	0.13	0.08	-0.07	-0.55	-1.38	0.00	-2.02	1.05
GBP HW J	-1.33	0.73	-0.46	0.13	-0.14	-0.09	-1.60	-9.95	20.05
AUD HW J	-0.62	0.05	-1.40	0.64	-0.47	-0.42	3.08	-7.65	7.58
Eur MS Cts	-0.36	-1.11	0.44	0.44	-0.54	1.51	5.59	1.68	3.38
JPY MS Cts	-1.18	-0.93	0.86	-0.57	0.18	5.70	4.43	5.02	3.48
GBP MS Cts	-1.06	0.30	0.04	0.30	-0.12	1.10	0.26	-5.73	15.31
AUD MS Cts	0.08	0.03	0.03	-0.53	-0.67	1.58	-3.17	8.02	-0.89
Eur MS J	-1.10	0.70	-0.05	0.07	0.01	0.08	-1.36	-10.13	21.02
JPY MS J	-0.71	0.27	0.19	-0.76	0.40	1.24	1.37	1.67	0.99
GBP MS J	-0.95	0.52	0.12	0.17	0.04	0.04	-0.87	-6.90	14.89
AUD MS J	-0.01	0.22	0.81	0.00	-0.70	1.26	2.52	-4.73	5.33

Note: The table shows t-statistics from the regression of Shapley-Owen sensitivity series on macro and monetary surprises (equation 4.4). Positive statistics denote events that tend to weaken the given statistical pattern while negative statistics denote events that strengthen the estimated pattern. $|t - statistics| > 1.96$ are highlighted.

Figure 1: R^2s by series and intraday segment

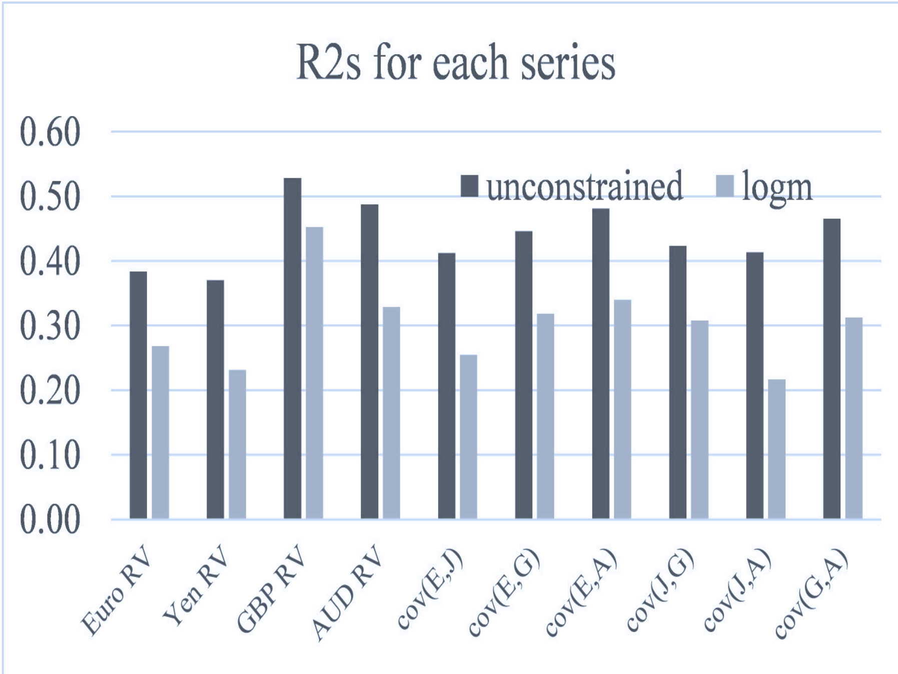

(a)

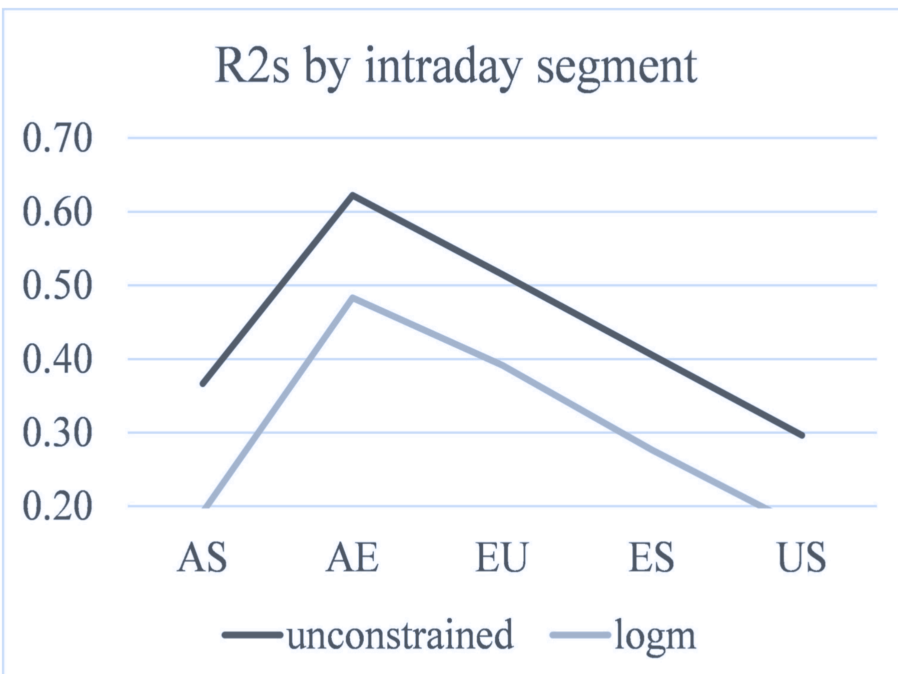

(b)

Note: The upper panel compares average R^2s over the intraday segments for each of the 10 series by method of estimation, unconstrained linear and logm models. The lower panel shows average R^2s over the 10 series for each of the 5 segments by method of estimation.

Figure 2: Shapley-Owen ratios by estimation method and intraday segment

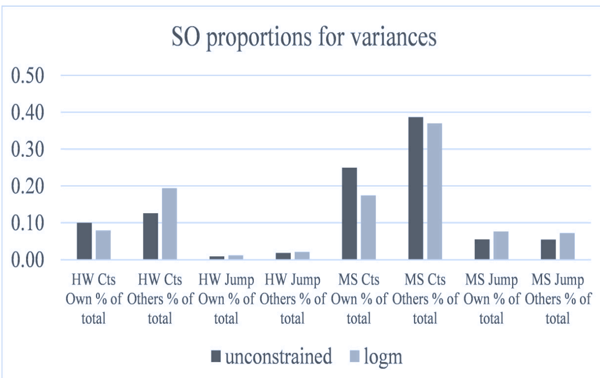

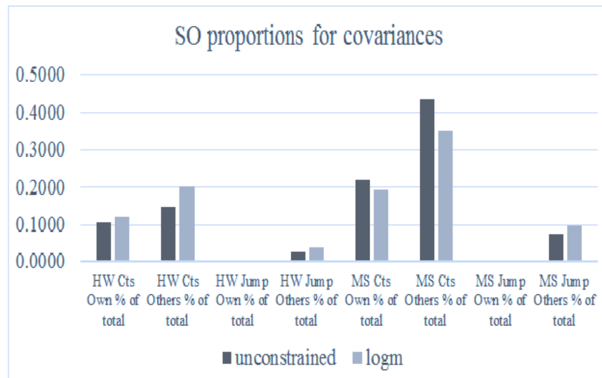

(a)

(b)

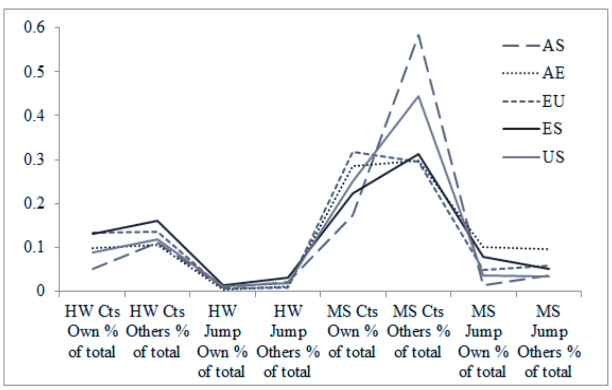

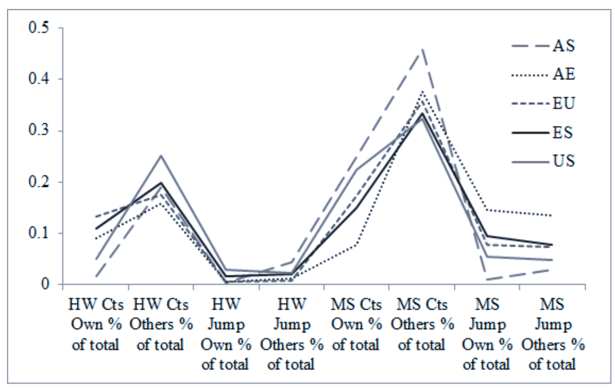

(c)

(d)

Note: The upper panels compares SO ratios for each of the 8 groups, averaged over the 4 variance (left) and 6 covariance (right) series, respectively. The SO ratios show the average proportion of the total R^2 that each group of explanatory variables accounts for under the Shapley-Owen metric. For example, a SO ratio of 0.10 means that the given group would account for 10% of the total predictability for the given set of dependent variables. The lower panels show the intraday variability of the 8 SO ratios implied by the unconstrained linear (left) and *logm* (right) models.