Appendix

A. The efficiency gain associated with intraday volatility measures

To illustrate the potential efficiency gains associated with the intraday return variability measure in equation (2) relative to the standard measure in equation (1), consider the extreme case where volatility remains constant within each day; i.e., $F_J \equiv F_{[J]}$ for $J \geq 0$, and [·] denotes the integer value operator. Given the distributional assumptions regarding P_J, it follows that

$$E\left(|R_t|\right) = \hat{F}_{t-1} \cdot (2/B)^{1/2} \quad , \tag{A1}$$

and

$$E\left(\sum_{n=1}^{N} |R_{t,n}| \right) = N^{1/2} \cdot \hat{F}_{t-1} \cdot (2/B)^{1/2} \, , \tag{A2}$$

which suggest the following two ex-post measures of the daily volatility

$$\hat{F}_{t,1} = (B/2)^{1/2} \cdot |R_t| \quad , \tag{A3}$$

and

$$\hat{F}_{t,N} = N^{-1/2} \cdot (B/2)^{1/2} \cdot \sum_{n=1}^{N} |R_{t,n}| \quad . \tag{A4}$$

While both estimators are unbiased, the latter is vastly superior. Specifically,

$$\mathrm{Var}(\hat{F}_{t,N}) = N^{-1} \cdot (B/2) \cdot \mathrm{Var}\left(\sum_{n=1}^{N} |R_{t,n}| \right)$$

Thus, with $N = 288$ five-minute intraday returns, the standard deviation is

$$= N^{-1} \cdot (B/2) \cdot \left[N \cdot \mathrm{Var}(|R_{t,n}|) + 2 \cdot \sum_{i=1}^{N-1} (N-i) \cdot \mathrm{Cov}(|R_{t,n}|, |R_{t,n-i}| \right.$$

$$= (B/2) \cdot \mathrm{Var}(|R_{t,n}|) = (B/2) \cdot N^{-1} \cdot \mathrm{Var}(|R_t|)$$

$$= N^{-1} \cdot \mathrm{Var}(\hat{F}_{t,1}) \; . \tag{A5}$$

reduced by a factor of almost seventeen. While the intraday volatility dynamics are much more complex than assumed above, the calculation is suggestive of the greatly improved ex-post measurement of the latent volatility process afforded by the cumulative absolute returns. For a theoretical exposition on related issues, see Nelson (1992), and Nelson and Foster (1995, 1996).

B. Consistency and robustness of the flexible Fourier form regression

The statistical properties of the FFF-regression is determined by the properties of the error process, $\hat{u}_{t,n}$, in equation (6). It takes the following form

$$\hat{u}_{t,n} = \left(\log s_{t,n}^2 - E[\log s_{t,n}^2]\right) + \left(\log F_{t,n}^2 - \log \hat{F}_{t,n}^2 - E[\right. \\ + \left(\log Z_{t,n}^2 - E[\log Z_{t,n}^2]\right) . \qquad \textbf{(A6)}$$

This error term consists of three component processes. The last is simple, as it constitutes an i.i.d. process. The first captures the discrepancy of calendar and announcement components from their expected values. Such divergences arise from stochastic components in the intraday "seasonal" or "news" innovations that differ from their expected values. As such, errors from this source are the rule rather than the exception. Nonetheless, if the mean effects are correctly specified and the errors are stationary, this does not affect the consistency of the OLS estimator. The second term reflects potential misspecification of the estimated volatility component, $\hat{F}_{t,n}$. Given the complexity of this process, it is inevitable that any preliminary estimator is misspecified, so this error term is likely heteroskedastic, serially correlated and perhaps even biased. However, any bias is absorbed in the constant, *c*, and will not further affect inference. Moreover, as long as the regressand and the volatility process itself are stationary, this entire error component is stationary. We conclude that the OLS estimator is consistent, while the associated error process will display dependencies of unknown form. Consequently, formal inference requires the use of robust standard errors that are consistent under general heteroskedasticity and autocorrelation.

Another important issue is the robustness of the FFF-regression to outliers. Simple diagnostics point to a potentially serious problem, as the kurtosis of the five-minute return series is 21.5 compared to 4.5 for the 12-hour returns. However, the problem is effectively eliminated by the log transform. In fact, inspection of the (transformed) regressor series, $\$_{t,n}$, now suggests a possible "inlier" problem, arising from the low values obtained when taking logs of small positive squared returns. The problem is similar to that encountered when applying the Kalman filter to log-squared returns in order to estimate stochastic volatility models, see, e.g., Harvey, Ruiz and Shephard (1994). However, we explicitly analyze the data for the presence of unduly influential observations, following the procedure in Davidson and MacKinnon (1993), section 1.6. We also truncate the observations for $\$_{t,n}$ from below by letting all return observations in the interval (0% , 0.00036%) equal 0 percent (minus the sample mean) before transforming to $\$_{t,n}$. It was confirmed in both cases that the presence of inliers did not exert an appreciable impact on the estimated volatility pattern.

C.1 Regional trading segments, holidays, and data gaps

We begin by formally defining the regional trading segments. This classification is used to assign dummy variables to the intervals affected by regional holidays. The observance of Daylight Savings Time in Europe and North America, at periods that do not fully coincide, induce us to operate with four separate categories. Furthermore, the classification is not exhaustive, in the sense that there are periods which do not belong to any specific regional segment. This is immaterial since it is only used to specify periods that are affected significantly by regional holidays in one of the market centers. The following daily Greenwich Mean Time (GMT) trading zone definitions are used:

	Year Round
Wellington (New Zealand):	20:55-22:00
Sydney (Australia):	20:55-00:00
Tokyo (Japan):	00:00-06:00

	09/27-10/23	**10/26-03/26**	**03/26-04/02**	**04/05-09/24**
London (Europe):	07:00-15:00	07:00-16:00	06:00-15:00	06:00-15:00
Europe-N.America Overlap:	11:30-15:00	12:30-16:00	12:30-15:00	11:30-15:00

New York (N.America): 11:30-20:30 12:30-20:30 12:30-20:30 11:30-20:30

Regional holidays affect the entire trading segment, except for certain minor U.S. holidays, where an appreciable drop in quoting and trading activity only takes place after the London market closes. The following holiday periods were identified from the quote intensity as well as the Reuter's news tape.

	Dates	**Time Period**	**Occasion**
United States	10/12	11:30-20:30	Columbus Day
	11/11	16:00-20:30	Veterans Day
	11/26	12:30-20:30	Thanksgiving
	12/21-01/01	All Day	Christmas/New Year
	01/18	16:00-20:30	King's Birthday
	02/15	12:30-20:30	President's Day
	04/08	15:00-20:55	Easter Begins
	04/09	All Day	Easter
	04/12	20:55-20:30	Easter Ends
	05/30	11:30-20:30	Memorial Day
	07/05	11:30-20:30	July 4
	09/06	11:30-20:30	Labor Day

Tokyo - Dates: 11/03, 11/23, 01/15, 02/11, 04/29, 05/03, 09/15, 09/23

Wellington - Dates: 10/26, 01/25, 06/07

Sydney - Dates: 10/05, 01/26, 04/26, 06/14

London - Dates: 05/03, 05/30, 08/30

We also checked for slowdowns associated with regional holidays in a number of additional countries, including Hong Kong, Taiwan, Singapore, Germany, and Switzerland, but no clear signs of an effect could be detected, so these holidays were not included in the analysis.

All five-minute intervals, covered by the holiday periods listed above, were assigned one of two different dummies. The "Holiday"-dummy refers to periods of reduced activity, where reliable returns may nonetheless be obtained. An interpretation is that this corresponds to lower levels of general economic activity, where less relevant economic news are generated. The "Market Closure"-dummy refers to periods where the quoting intensity is so low as to render return calculations unreliable. Among the above holidays, the following are allocated to the latter "Market Closure" category:

	Dates	Time Period	Occasion
Market Closures:	10/12	15:00-20:30	Columbus Day
	11/26	16:00-20:30	Thanksgiving
	12/22	20:30-20:55	Christmas
	12/23-12/25	All Day	Christmas
	12/28	21:00-23:00	Christmas
	12/31	17:00-20:55	New Year
	01/01	All Day	New Year
	02/15	16:00-20:30	President's Day
	04/08	20:30-20:55	Easter
	04/09	All Day	Easter
	04/12	20:55-20:30	Easter
	05/30	06:00-20:30	Memorial Day
	07/05	11:30-20:30	July 4
	09/06	11:30-20:30	Labor Day

The trading restrictions in Japan over the sample period precludes reliable assessment of the properties of the return series over the local lunch period. It effectively corresponds to a "weekend" in the midst of the trading day. Formally, we define a market closure each day during

Tokyo Lunch-Time: 03:00-04:45

Finally, we identified some apparent failures in the data transmission which result in lengthy gaps in the quote series. All of the affected intervals were treated as market closures. The specific periods are:

	Dates	Time Period
Data Gaps:	10/21	01:18-05:37
	10/28-29	22:16-01:15
	11/17	01:30-05:39
	12/16	01:15-05:12
	01/08	00:33-06:20
	02/10	01:35-06:27
	02/22	04:52-06:40
	05/21	16:41-21:00
	09/26-27	21:57-06:07

The market closures present a modeling dilemma, since we want to eliminate these observations, but also want to retain the strict periodicity associated with the intradaily and weekly features of the high frequency return series. We solve this by artificially assigning a very low, positive return (standardized by an overall daily volatility factor) to all these intervals, and then removing (zeroing out) all regressors except the market closure-dummy from these intervals. This implies that the dummy "explains" the low returns (near) perfectly, while the inference regarding all other features of the return series is unaffected.

C.2 Constrained calendar and announcement volatility response patterns

In order to accommodate the overall impact through a parsimonious representation that also allows for efficient inference, the reported estimates for the announcement and calendar effects are based on the imposition of an a priori structure on the volatility response pattern. In particular, assuming that the feature in question affects volatility from interval n_0 to $n_0 + n_1$, the impact over the event window, $J = 0, 1, \ldots, n_1$, may then be represented by a polynomial specification,

$$p(J) = c_0 + c_1 \cdot J + \ldots + c_P \cdot J^P. \qquad \textbf{(A7)}$$

Of course, for $P = n_1$ this would effectively imply the estimation of a dummy variable for each of the $\bar{N} / n_1 + 1$ event intervals. However, the use of a lower order polynomial affords a great degree of flexibility along with a significant reduction in the dimensionality of the parameter space. Furthermore, sensible constraints on the response pattern, including smoothness, are readily imposed in terms of the polynomial representation. For example, the requirement that the impact reflects a gradual movement away from the standard pattern is imposed by enforcing $p(0) = 0$. This simply annihilates the constant, i.e., $c_0 = 0$. Another desired property may be that the effect slowly fades, which is obtained by imposing $p(\bar{N}) = 0$. Substituting $J = \bar{N}$ into $p(J)$, solving for c_P, and inserting the resulting expression for c_P back into $p(J)$, leads to a restricted polynomial with one less parameter,

$$p(J) = c_0 \cdot [\, 1 - (J / \bar{N})^P \,] + c_1 \cdot [\, 1 - (J / \bar{N})^{P-1} \,] \cdot J + \ldots + c_{P-1} \cdot [\, 1 - (J / \bar{N}) \,] \cdot J^{P-1}. \qquad \textbf{(A8)}$$

We can now classify a number of our calendar and all of our announcement regressors through the choice of polynomial order, *P*, the response horizon, or $\bar{N}$, and the endpoint constraints imposed on *p(J)*. The following specifications underlie the results reported in the paper:

Tokyo Market Opening:	$\bar{N} = 6$,	$P = 1$,	$p(\bar{N}+1) = 0$,
Late Summer Day Slowdown:	$\bar{N} = 60$,	$P = 2$,	$p(0) = p(\bar{N}+1) = 0$,
Early Monday Effect:	$\bar{N} = 17$,	$P = 2$,	$p(\bar{N}+1) = 0$,
Late Friday Effect:	$\bar{N} = 46\ (58)$,	$P = 2$,	$p(0) = 0$,
EMS-Band Widening:	$\bar{N} = 30$,	$P = 3$,	$p(\bar{N}+1) = 0$,
Employment Report:	$\bar{N} = 24$,	$P = 3$,	$p(\bar{N}+1) = 0$,
All Other Announcements:	$\bar{N} = 12$,	$P = 3$,	$p(\bar{N}+1) = 0$.

The above representations leave one free parameter for the Tokyo market opening and the Summer slowdown, and two free parameters for the weekend effects denoted "Monday early" and "Friday late". The "Friday late" coefficients are identical in Summer and Winter, but the effects lasts an additional hour during Summer due to Daylight Savings Time. Finally, there are three announcement effect parameters, but as explained in Section IV.A, we further restrict this pattern by imposing the common structure,

$$p_k(J) = 8_k \cdot p_0(J)\ , \tag{A9}$$

where $p_k(J)$ refers to the polynomial for event type *k*, and $p_0(J)$ denotes a fixed response pattern. Specifically, we calibrate the pattern by fitting all three parameters for a set of announcements of about equal significance, resulting in a benchmark pattern that resembles the one associated with Category I releases. Concretely, (c_0, c_1, c_2) = (2.18868, -0.64101, 0.07663). This uniquely identifies $p_0(J)$, and $p_k(J)$ thus has only one free "loading" parameter, 8_k. Of course, this procedure only strictly applies for response horizons corresponding to $\bar{N}$ = 12. In order to retain the benchmark pattern for larger $\bar{N}$, we let the *J*-variable progress only by a (12/$\bar{N}$)-fraction of a unit per five-minute interval, rather than a full interval. This "stretches" the event time scale so that it conforms to the desired horizon.

Finally, we apply the corresponding "time-deformation" procedure to the sinusoids in the U.S. Summer Time intraday pattern in order to compensate for the one hour leftward shift from 7:00 to 6:00 GMT. This elongation of the intraday pattern is implemented over 19:55 to 00:00 GMT.