

How Is Macro News Transmitted to Exchange Rates?

Martin D. D. Evans*

Richard K. Lyons

31 July 2003

Abstract

This paper tests whether macroeconomic news is transmitted to exchange rates via induced transactions, and if so, what share occurs via transactions versus traditional direct adjustment of price. We identify the link between order flow and macro news using a heteroskedasticity-based approach. This involves jointly testing (1) whether macro news flow increases order flow volatility and (2) whether the induced order flow has signed (first moment) effects on the exchange rate. The answer to both questions is yes: in both daily and intra-daily data, order flow is considerably more volatile when macro news is flowing, and these signed orders have the theoretically predicted effects on the exchange rate's direction. Of news' total price effect, induced order flow accounts for two-thirds, with direct news effects accounting for one-third. In terms of total exchange rate variation, the order flow channel brings news' explanatory power up to 30 percent, versus estimates in the 1–5 percent range from existing literature, helping to resolve the puzzle of missing news effects.

Correspondence

Richard K. Lyons

Haas School of Business

U.C. Berkeley

Berkeley, CA 94720-1900

Tel: 510-642-1059, Fax: 510-643-1420

lyons@haas.berkeley.edu

faculty.haas.berkeley.edu/lyons

* Respective affiliations are Georgetown University and NBER, and UC Berkeley and NBER. This paper was circulated under the previous title "Why Order Flow Explains Exchange Rates." We thank the following for helpful comments: Brian Boyer, Anusha Chari, Frank Diebold, Jon Faust, Simon Gervais, Aditya Kaul, Antonio Mello, John Rogers, Andrew Rose, and seminar participants at the Empirical Exchange Rate Models conference (Sept. 2001), the American Finance Association meetings (Jan. 2002), the Yale Market Microstructure conference (June 2002), U.C. Berkeley, Alberta, Michigan, London Business School, the New York Fed, Aix-Marseille, Virginia, Wharton, Insead, LSE, the Norwegian central bank, and the Italian central bank. Both authors thank the National Science Foundation for financial assistance, which includes funding for a web clearinghouse for micro-based research on exchange rates (at faculty.haas.berkeley.edu/lyons).

How Is Macro News Transmitted to Exchange Rates?

“Combining our perspectives focuses attention on the causal links among news, order flow, and forex movements, which in our view is a prime candidate for future research.”

Andersen, Bollerslev, Diebold and Vega (2003, p. 59)

In macroeconomic models of exchange rates, news maps directly into prices. The effect of news on currency demands in these models is common knowledge and transactions—though perhaps engendered by the change in exchange rates—play no role in causing the change. In microeconomic models of asset prices, in contrast, transactions do play a causal role in price determination (e.g., Glosten and Milgrom 1985, Kyle 1985). The causal role arises because transactions convey information that is not common knowledge. In this paper we test whether any part of the effect of news on exchange rates is transmitted via transactions and, if so, what share is transmitted that way versus the traditional direct channel.

That transactions might play a role is motivated by recent empirical work demonstrating a link between signed transaction volume (order flow) and signed exchange rate changes.¹ In the models employed by these papers, order flow affects exchange rates as a proximate determinant. The underlying determinant, which theory labels information, is not specified, nor is it directly tested. This leaves open the nagging question of what really drives the order flow. This paper is an attempt to address that question by examining whether macroeconomic news might be a determinant of signed transaction volume. This question is distinct from whether unsigned volume is determined by news, a well-established property of many speculative markets (see, e.g., Fleming and Remolona 1999).²

¹ This evidence is from both micro (i.e., single marketmaker) and macro (marketwide) studies. See, e.g., Payne (1999), Rime (2000), Evans and Lyons (2002a) and Evans (2002). Order flow—a concept from microstructure finance—is the cumulation over time of *signed* trades. Trades are signed according to whether the initiator is buying or selling. (The marketmaker posting the quote is the non-initiating side.) For example, a sale of 1 unit by a trader acting on a marketmaker’s quote is order flow of -1 . A large empirical literature within finance shows that signing trades this way provides considerable explanatory power (see, e.g., the reviews in Madhavan 2000 and Lyons 2001).

² Idiosyncratic portfolio rebalancing in response to news is consistent with an effect on volume. However, under rational expectations, given that immediate adjustment in price is unbiased, on

We estimate a trading model that distinguishes three sources of exchange-rate variation. The first source mirrors traditional models—macro news that is impounded in price immediately and directly (i.e., with no role for transacted order flow). The second source is the indirect effect of news via induced order flow.³ The third source is order flow unrelated to public news (such as that due to banks' changing risk tolerances, firms' changing hedging demands, or individuals' changing liquidity demands; see, e.g., Evans and Lyons 2002a). We find that for the DM/\$ market, all three sources of exchange rate variation are significant. The flow of transacted orders between marketmakers (our measure of order flow) varies considerably with macro news flow, such that roughly two-thirds of the effect of macro news on exchange rates is transmitted via order flow, the remainder being the direct effect of news. This is consistent with our finding that, though order flow is always important for determining exchange rates, it is even more important when news is flowing than when it's not (contrary to intuition from macro models). With both channels operating, news accounts for about 30 percent of total price variance, which is an order of magnitude larger than findings from past work (more below).

Though the literature linking exchange rates to news is vast, we are the first to our knowledge to use order flow to sort out the relationship. The existing literature has two branches: the first attempts to explain the direction of exchange-rate changes (first moments) and the second, later branch attempts to explain exchange-rate volatility (second moments). A common finding of the first branch is that, at least at the daily frequency, directional effects from scheduled macro announcements are difficult to detect because they are swamped by other factors affecting price. Intraday event studies do find statistically significant effects, particularly for employment and money-supply announcements (Andersen et al. 2003).⁴ The second, later branch of this literature—which focuses on news effects on volatility—is partly

average good news for the dollar should not produce positive (or negative) order flow; i.e., one would not expect a relative increase in executed transactions initiated by dollar buyers. We return to this issue below when discussing potential simultaneity bias.

³ That this channel might be operative is another reminder that order flow and demand are not one to one. To clarify, consider a simple counter-example. In the common-knowledge tradition of exchange rate economics, when positive public news arrives, demand increases, causing price to increase—without any flow of transacted orders occurring or needing to occur. This is incompatible with demand and order flow being one to one: the demand shift does not generate any order flow at the new price.

a response to early difficulty in finding news effects on first moments.⁵ This work finds that arrival of scheduled announcements does indeed produce the largest exchange-rate changes. On the other hand, though major announcements dominate the volatility picture within short release windows, the ability of these fundamentals to account for volatility changes overall is lower than that of less fundamental features such as time-of-day effects and ARCH (Andersen and Bollerslev 1998).⁶

Our paper departs from earlier work on macro news in two main ways. First, we consider a broader set of macro news events. Scheduled announcements account for about 10 percent of the macro news arrivals we obtain from the Reuters Money Market Headline News screen. (The median number of macro news arrivals per day in our sample is 11.) This allows us to test whether this wider set of news types affects the processes for transactions and returns in ways not examined previously. At the same time, including a broad set of news types requires a change in estimation strategy because it is not possible to measure ex-ante expectations for most of them (e.g., a report on the Headline News screen of an official stating that the trade deficit is unsustainable). This leads to the second of our main departures from earlier work: like Rigobon and Sack (2002), we depart from the event-study approach and adopt instead an approach based on the state dependent heteroskedasticity.⁷ Specifically, we identify (via GMM) the relative importance of direct and indirect effects from news by allowing the variances of shocks to order flow and price to depend separately on the rate of news arrival. These second moment conditions identify our model and answer the central question before us, namely, whether order flow plays a role in transmitting news to exchange rates. This approach, as in Rigobon and Sack (2002), does not require measurement of news'

⁴ See also, for example, Cornell (1982), Engel and Frankel (1984), Hakkio and Pearce (1985), Ito and Roley (1987), Hardouvelis (1988), Klein (1991), and Ederington and Lee (1995).

⁵ See, for example, Goodhart et al. (1993), DeGennaro and Shrieves (1997), and Andersen and Bollerslev (1998). See also the work on bond prices and announcements, e.g., Fleming and Remolona (1999), Balduzzi et al. (2001), Fleming (2002), and Green (2002). The latter two papers are especially relevant in that they use direct measures of order flow in fixed income markets. Green (2002), for example, finds evidence that asymmetric information increases following public macro announcements.

⁶ Accounting for volatility changes and accounting volatility levels of are not the same. For example, abnormal volatility in response to heightened news arrivals may dissipate rapidly, but trading induced by past news arrival can continue to contribute to normal volatility (a point we return to below).

unanticipated component; it requires the weaker assumption that one can distinguish periods in which the variance of macro news flow is relatively high. We use news arrivals for this purpose.

Before delving into details, it is worthwhile recognizing the big picture into which this paper fits. The vast event-study literature noted above that addresses news effects on exchange rates has in many instances been quite successful in explaining signed exchange rate variation *within event windows*, but remains relatively unsuccessful in accounting for total variation. A good example at the daily frequency is Klein (1991), who regresses signed exchange rate changes on signed trade-balance news and finds that news can explain about 40 percent of exchange rate changes on those days. This is a lovely result. But remember that trade balance news arrives monthly. That means that roughly 95 percent of exchange rate variation is not included in the regression (20 of 21 trading days per month). Thus, an R^2 statistic of 40 percent is still accounting for less than 3 percent of total exchange rate variation. A good example from the intraday event-study literature is Andersen et al. (2003). They too find within-event-window R^2 's that are genuinely impressive. But as they note (p. 50), their five-minute post-event intervals represent only 0.2 percent of their sample observations (there just aren't that many scheduled macro announcements). Thus, even with event-window R^2 's of 50 percent, and with event-window variance increasing by a factor of 20, the total exchange rate variation being explained would be only 2 percent.⁸ In our paper, based on a much broader sample of macro news events, and allowing for an indirect order flow channel, we find that news' power to explain total exchange rate variation is an order of magnitude higher at roughly 30 percent.

⁷ See the discussion in Rigobon and Sack (2002) comparing the merits of the event-study and heteroskedasticity approaches. Omitted variable bias in event-study analysis is simply a manifestation of a point made above, namely, that event effects are often swamped by other factors affecting price.

⁸ That security return volatility is not constant over time was well documented by French and Roll (1986). Our daily-frequency example from Klein (1991) could include two adjustments in this respect: exchange rate volatility over weekends is not zero, which lowers his overall explanatory power, but announcement days tend to have higher volatility than non-announcement days, which raises his explanatory power. Neither of these adjustments is large enough to alter the basic message. Our intraday example from Andersen et al. (2003) is conservative in that the 50 percent R^2 is on the high side (see their Table 2) and the variance ratio of 20 is on the high side as well (even for the highest-impact announcement, the Employment Report, instantaneous variance increases by roughly a factor of 10; see Andersen and Bollerslev 1998, page 249).

The remainder of the paper is in four sections. Section 1 describes our data and presents some descriptive statistics. Section 2 presents daily analysis of the three sources of price adjustment. Section 3 presents intra-daily analysis of the three sources. Section 4 concludes.

1. Data and Descriptive Statistics

Our empirical strategy is two pronged, encompassing two data frequencies and two complementary approaches to how macro news is transmitted to exchange rates. The first approach applies at the daily frequency, the second at intraday frequencies (in our case, five-minute observations). Daily analysis of exchange rate variation is interesting because daily changes in nominal exchange rate are, to a first approximation, a martingale. Any empirical model that explains daily price increments is therefore relevant for explaining exchange rate levels at long horizons (i.e., one cannot sensibly argue that these are rapidly dissipating price effects). This martingale property at the daily frequency does not apply to intraday prices, which exhibit mean reversion, as we shall show. The upshot is that disentangling different sources of daily exchange rate variation provides a solid indication of decompositions at frequencies more familiar to macroeconomists (e.g., monthly). But daily analysis does not convey the rich intraday structure of news effects. Nor is it suited to addressing whether news makes order flow more or less important in exchange rate determination. For these we turn to intraday analysis.

Our order flow and price data are drawn from time-stamped, tick-by-tick transactions in the largest spot market—DM/\$—over a four-month period, May 1 to August 31, 1996 (full 24-hour trading day). The data source is the same as that in Evans (2002) and we refer readers there for additional detail. The transactions are from the Reuters Dealing 2000-1 system. At the time of our sample, Dealing 2000-1 was the most widely used electronic dealing system: according to Reuters, over 90 percent of the world's bilateral transactions between DM/\$ marketmakers took place through the system.⁹ For every trade executed on D2000-1, our data set

⁹ At the time of our sample, transactions between marketmakers accounted for about 75 percent of total trading in major spot markets. This 75 percent breaks into two transaction types—direct (bilateral) and brokered. Direct trading accounted for about 60 percent of trades between marketmak-

includes a time-stamped record of the transaction price and a bought/sold indicator. The bought/sold indicator allows us to sign trades for measuring order flow.

For our daily analysis, Δp_t is the change in the DM transaction price for purchases of dollars from 5 pm BST (British Summer Time) on day t-1 to 5 pm BST on day t.¹⁰ When a purchase transaction does not occur precisely at 5 pm, we use the immediately preceding purchase price. Daily order flow x_t is the difference between the number of trades initiated by dollar buyers and the number initiated by dollar sellers over the same time interval (in hundred thousands, negative sign denotes net dollar sales).¹¹ For our five-minute returns, which we denote Δp_i , we also use the immediately preceding transaction when no transaction occurs precisely at the five minute mark. (With roughly 1 million transactions per day, the preceding transaction is generally only seconds earlier.) Five-minute order flow x_i is defined similarly to x_t but over the five-minute interval matching Δp_i .

Our news data come from the Reuter's Money Market Headline News screen (source: Olsen Associates). These screens are standard equipment on FX trading desks and are used for high frequency monitoring by non-marketmaker participants as well. We cleaned the raw news data by excluding news items of the following four types: (i) announcements of upcoming known holidays, (ii) announcements that a scheduled release would take place (e.g., "Monthly employment report due out tomorrow"), (iii) duplicate announcements (the same news is repeated with a slight change in wording), and (iv) announcements referring to the dollar/DM exchange rate. The first three of these filters are intended to distill information that is truly incremental. The fourth is intended to protect against feedback from exchange rates to macro news flow (e.g., an announcement that the dollar/DM exchange rate had been volatile that morning would be a clear problem for interpreting macro news

ers and brokered trading accounted for about 40 percent. For more detail on the Reuters Dealing 2000-1 System see Lyons (2001) and Evans (2002).

¹⁰ Using prices from purchase transactions (i.e., transactions at the ask) eliminates return reversals that would arise in prices that bounce randomly from bid to ask.

¹¹ In direct trading between marketmakers, orders sizes are standardized, so variation in size is much smaller than variation in the size of individual customer-marketmaker trades. Note too that using measures of order flow based on numbers of transactions rather than size is common in work on equity markets, even when both measures are available (see, e.g., Hasbrouck 1991). Our data set does include total dollar volume over our sample, which allows us to calculate an average trade size; we use this below to interpret the estimated coefficients.

flow as exogenous; there weren't, however, any actual announcements with such clear simultaneity with exchange rates). In all, the four filters excluded less than 10 percent of news arrivals.

Our daily frequency variable A_t is the number of news arrivals relating to U.S. or German macroeconomics between 5:01 pm BST on day $t-1$ and 5 pm BST on day t . The five-minute analogue of A_t is denoted A_{5t} . The dummy variable A_{5t} takes the value of one if there was a news arrival (U.S. or German) during the previous five minutes. (We use a dummy variable approach in the five-minute data because there are very few instances of more than one news arrival during a single five-minute observation window.) In our four-month sample, there are 315 five-minute windows in which at least one news arrival occurs. Arrivals in the raw news data recorded after 20:00 hrs and before 6:30 (BST) are quite rare and typically refer to previous news events within the heart of the trading day. For this reason we do not include the few news arrivals that occur in this time interval in our sample.

As noted, our methodology allows us to consider a much wider set of macro news types.¹² This is important because our four-month sample sharply constrains our ability to work with news arrivals on a fully disaggregated basis (i.e., by individual news types, such as Unemployment Claims). Though there is no ready measure of ex-ante expectations for more than 80 percent of this full news set, there is a subset for which a commonly used measure of expectations exists, namely the subset of scheduled macro announcements (that measure being survey responses provided by Money Market Services and used extensively in foreign exchange event studies). We obtained these survey data from Money Market Services so that we could measure innovations for this subset more accurately and compare our results from this subset to results from the full set of macro news events. This provides an indication of the statistical power gained by extending the news definition beyond the narrow definition for which the best expectation proxies exist. Another (partial) disaggregation for which our sample provides some statistical power is U.S. versus

¹² A wider set of news types is also considered by Eddelbuttel and McCurdy (1998), though their focus is the effect of news flow on exchange rate volatility, not the joint dynamics of news, signed order flow, and signed price movements. Berry and Howe (1994) use the number of public releases as a news flow measure to address equity market volume and volatility.

German news. Results for these two disaggregations are presented in the next section.

Table 1 presents descriptive statistics for the variables used in our daily analysis (upper panel) and intraday analysis (lower panel). For the daily analysis these variables include news flow A_t , price changes Δp_t , and order flow between marketmakers x_t . The median number of daily news arrivals is 11, 8 of which pertain to Germany. As noted, these news events extend well beyond the scheduled announcements that are the focus of most past work. (The most important scheduled announcements for explaining changes in exchange rate volatility are employment, GDP, the trade balance, and durable goods; see Andersen and Bollerslev 1998.) Our intraday analysis includes a variable not included in the daily analysis, namely n_i , the total number of trades over each five-minute interval. This variable accounts for dependence on trading intensity in the intraday model, as we shall see. Note that intraday price changes display (negative) autocorrelation, but only at lag one. Thus, future price changes cannot be predicted from past changes beyond a five minute horizon—a result confirmed by the ARMA model estimates in Evans 2002 (i.e., price changes follow an MA(1)). Order flow, at least at the five-minute frequency, is much more persistent. (Order flow exhibits virtually no persistence at the daily frequency.) For perspective, we note that on the day with the most news arrivals (28), the realization of $|\Delta p_t|$ was at the 73rd percentile of its distribution, the realization of $|\Delta x_t|$ was at its 92nd percentile, and the realization of trading intensity n_i was at its 86th percentile.

Figures 1 and 2 provide more perspective on the announcement data. Figure 1 shows the sample distribution for all news arrivals (both U.S. and German) by 30-minute interval. Figure 2 provides a list of all the arrivals in the first three days of the sample (19 in total). These news events clearly extend well beyond scheduled announcements.¹³ (Including only the latter would leave many trading days in our sample without a single news event.) Are these news arrivals just noise? We prefer

¹³ Our filters did not exclude the third news arrival because it is an interpretation of the second, not simply a restatement. The filters did not exclude the eighth arrival because in Germany (unlike the U.S.) the timing of this type of release is not regular. Arrival of this type of information can certainly affect the timing of participants' position taking.

to let the data speak: if they are just noise then their arrival rate should not be correlated with trading activity or exchange rate volatility. We will test this. We would also stress that in any event, data noise is much more significant problem for a paper whose contribution is a negative result, since that negative result could stem from the noise. The findings in this paper, in contrast, are positive.

Whether news arrival is correlated with trading activity and volatility is addressed in Figure 3. It presents a scatter plot of the number of news arrivals per day (A_t) against the daily variance of returns and order flow. (To measure the two realized variances more precisely, we use a measure of integrated variance based on a five-minute sampling frequency.¹⁴) From the figure there is clearly a relationship: the variances of both returns and order flow are higher when the flow of macro news is higher. The positive link between the variance of returns and news flow is well documented in the existing literature. The positive link between the variance of order flow and news flow is new. Our daily model gives this positive link more structure, to which we turn in the next section.

As a final point regarding data, the major spot markets have recently undergone some institutional changes. Since the time of our sample in 1996, trading between marketmakers has migrated significantly from bilateral trading (the source of our data) to brokered trading over the EBS and Reuters Dealing 2000-2 systems (see BIS 2002). Does this render our analysis based on order flow obsolete? No, for two reasons. First, bilateral trading between marketmakers accounts for the lion's share of total trading through virtually all of the floating-rate era (i.e., post Bretton Woods, which collapsed in the early 1970s). That the FX market is now more oriented toward brokered trading does not render the type of data we use here obsolete for understanding post Bretton Woods floating rates. Second, though it is true that the specific order flow data we use is now a smaller share of the total, this does not imply that order flow analysis is less useful going forward. Strategies for order flow measurement and capture will need to change (e.g., extracting it from electronic broker systems). But the role of order flow in conveying information, in theory and in practice, transcends market structure. The types of information that order flow conveys—particularly types with persistent price effects—are unlikely to

¹⁴ See the Appendix for details on these integrated variance measures.

change radically as the FX market changes over time. Put another way, the underlying information structure of this market has more to do with the properties of the asset being traded than it does with market structure per se. Order flow should continue to convey dispersed information that needs to be aggregated and subsequently impounded in price. The methodology we introduce here is more robust than a cursory look at market structure might suggest.

2. Daily Analysis: Direct vs Indirect Effects from News

Our daily analysis relies heavily on the portfolio shifts model of trading developed in Evans and Lyons (2002a). Accordingly, rather than review that model's details, our presentation focuses on ways this paper extends that earlier framework. We base our daily analysis on the portfolio shifts model because it serves three important purposes. First, the model clarifies how and why causation can run strictly from order flow to price. Within microstructure theory more broadly, this direction of causality is the norm (e.g., Glosten and Milgrom 1985 and Kyle 1985). Second, the model clarifies why price determination should depend most directly on a particular type of order flow, namely, the flow of orders between marketmakers. (This measure of flow is the most transparent to the marketmakers and therefore plays the most important role in updating their beliefs.) This feature of the model corresponds well to reality and produces estimable equations because this type of flow data is available (largely because this type of trading occurs electronically, unlike data on trades between marketmakers and the public). Third, the earlier version of the portfolio shifts model provides underlying economics for why price effects from order flow should persist. The basic reason is that flow innovations provide signals about stochastic portfolio shifts (such as those due to changing hedging demands by a subset of agents). These portfolio shifts must be reabsorbed in equilibrium by the rest of the market, so if the shifts are persistent, the effects on price are persistent too.

This paper's extension of the portfolio shifts model introduces macro news arrivals whose implications for the exchange rate are not common knowledge.¹⁵

¹⁵ An earlier version of this paper provides formal treatment of this change in information structure. Though the change is fundamental, it does not involve a fundamental change in the analytics, so we

Specifically, the information in news has two components. The first component is a common-knowledge (or “mean”) part: all agents agree what the first part’s impact on the exchange rate should be. (This will be the source of the direct effects on the exchange rate.) The second component is the part whose implication for the exchange rate is not common knowledge. Suppose, for example, that all agents do not have access to the same technology for transforming the macro data into an exchange rate forecast. The information in this second part is thus dispersed throughout the economy and for price-setting marketmakers to aggregate it, they need to learn from the public trades that are induced by it.¹⁶

The model produces three basic channels through which information affects prices. The first is the traditional direct channel for public information: the common-knowledge part of macro news is reflected instantaneously in price, a process that does not rely in any way on the flow of transacted orders. The second channel is for the part of news that is not common knowledge: information in the induced order flow is impounded in price as the flow is realized. The third is the channel for information in other types of order flow, i.e., order flow unrelated to macro news flow (such as trades motivated by shocks to individuals’ hedging demands or shocks to banks’ effective risk aversions due to capital constraints, etc.).

The model leads to the following empirical specification for daily price changes and order flow:

$$\Delta p_t = \alpha x_t + \xi_t + \kappa_t \tag{1}$$

$$x_t = e_t + \eta_t \tag{2}$$

where Δp_t is the change in the exchange rate (DM/\$) from the end day t-1 to end day t and x_t is order flow between marketmakers realized over the same period. The

omit it here to conserve space. A copy of that earlier version is available at faculty.haas.berkeley.edu/lyons/macronewsold.pdf.

¹⁶ How do marketmakers respond to scheduled announcements in practice? For the major scheduled announcements (which represent a small subset of our data), bank economists typically produce a forecast of the announcement’s value (e.g., the employment report) that is distributed to that bank’s marketmakers that morning. The sheet with the forecast will often include a prediction for the exchange rate impact, but marketmakers are left to their own judgment regarding how to respond. Our best reading from observing the process is that some price adjustment does indeed occur rapidly across the market, with little apparent role for flow. But announcements often induce follow-on trading, some

parameter α captures the price impact of order flow, i.e., it reflects the information content of order flow. Prices and order flow are subject to four shocks representing different sources of information hitting the market. These shocks are mean zero, mutually uncorrelated, and serially uncorrelated. The ξ_t and κ_t shocks represent information that is impounded in price directly. ξ_t is the common knowledge effect of macro news arrivals on the exchange rate. κ_t represents other factors directly impounded in prices, i.e., factors unrelated to both order flow or macro news events (possibly noise). Order flow is driven by the e_t and η_t shocks. The e_t shocks represent order flow effects from macro news arrivals—the non-common-knowledge effect of the news. Shocks to order flow that are unrelated to macro news (i.e., portfolio shifts arising from other sources such as changing risk tolerances or hedging) are represented by the η_t shocks.

Does it make sense that the common-knowledge effect of macro news arrivals ξ_t should be orthogonal to the order flow innovations e_t and η_t ? This is clearly important for identifying our model, so let us address it in more detail. That public information—understood to have common-knowledge implications—does not affect order flow is a working premise with a long history in empirical finance, dating back at least to the work of Hasbrouck (1991), and serving as the basis for much important work by various authors since then (see, e.g., Madhavan, Richardson, and Roomans 1997 and the survey in Madhavan 2000). To understand why it makes sense, it is important to distinguish clearly between (signed) order flow and (unsigned) volume. It is quite plausible that public news should induce trade in assets. But trade in assets is volume, not order flow. In general, the signed order flow implications of common-knowledge price adjustments are zero: the new price, if it is a market-clearing price, should not systematically favor imbalances of sell orders over buy orders, or vice versa (e.g., there should not be a correlation between bad public news for the DM and subsequent net DM sell orders, so long as the update of the market price is rational).

of which comes later the same day from customers whose trading response is not instantaneous. It is these (unforecastable) customer trading responses that the model is designed to capture.

The estimation strategy in our daily analysis relies on identification through conditional heteroskedasticity. The key conditioning variable is the number of macro news arrivals between the end of day $t-1$ and the end of day t , A_t . We allow the variance of the ξ_t and e_t shocks to increase with the number of news arrivals:

$$\text{Var}(e_t) = \sigma(A_t) \quad \text{and} \quad \text{Var}(\xi_t) = \omega(A_t) \quad (3)$$

where $\sigma(0) = \omega(0) = 0$, with $\sigma'(\cdot) \geq 0$ and $\omega'(\cdot) \geq 0$. This specification implies that neither the ξ_t nor the e_t shocks affect prices on days when there are no news arrivals. The shocks κ_t and η_t are independent of news, so their variances are unrelated to A_t . We assume these variances are constant:

$$\text{Var}(\kappa_t) = s_\kappa \quad \text{and} \quad \text{Var}(\eta_t) = s_\eta . \quad (4)$$

To estimate the model described in equations (1) – (4) we specify linear forms for the variance functions $\sigma(A_t) = \sigma A_t$ and $\omega(A_t) = \omega A_t$, which we later subject to diagnostic testing. The parameters of these functions along with s_κ and s_η are estimated by the Generalized Method of Moments (GMM), with conditions on second moments playing the central role in identification. (Estimation details are described in the Appendix.) The answer to the central question of this paper—does order flow help to transmit macro news to exchange rates—amounts to a test of whether the parameters σ and α are positive and significant (the latter being an estimate of the signed impact of order flow on prices). The relative importance of the direct and indirect news channels can be determined from the relative sizes of the estimated parameters ω and σ .

To the above specification we add one last feature: room for contemporaneous returns to enter the order flow equation (2). An effect of returns on order flow is not optimal within the motivating theory, but as an empirical matter, estimating an overly restrictive model is not necessary. Specifically, we enrich the model in equations (1) and (2) to the following:

$$\Delta p_t = \alpha x_t^* + \xi_t + \kappa_t \quad (5)$$

$$x_t = x_t^* + x_t^+ \quad (6)$$

$$x_t^* = e_t + \eta_t \quad (7)$$

$$x_t^+ = \phi \Delta p_t \quad (8)$$

This specification includes an additional component of order flow, x_t^+ , which is driven by contemporaneous returns. Though as econometricians we observe only the composite order flow x_t , the parameter ϕ governing x_t^+ is easily estimated within our GMM framework.

To illustrate how identification through heteroskedasticity is achieved in the model, consider how the two key parameters α and ϕ relate to the covariance between price changes and order flow. Equations (5)-(8) imply the following expression for that covariance:

$$\text{Cov}(\Delta p, x) = \alpha(1+\alpha\phi)\text{Var}(e+\eta) + \phi\text{Var}(\xi+\kappa).$$

Identification is achieved by considering how the moments in the data change across states. In particular, the arrival of macro news affects both terms on the right-hand side: from equation (3), both $\text{Var}(e)$ and $\text{Var}(\xi)$ increase, where the former is the indirect order-flow effect on price and the latter is the direct effect on price. The resulting change in $\text{Cov}(\Delta p, x)$ depends on the values of α and ϕ . While this link alone is insufficient to identify α and ϕ , these parameters also determine how macro news affects the variance of price changes and order flow (see appendix). Taken together, data on changes in $\text{Cov}(\Delta p, x)$, $\text{Var}(\Delta p)$, and $\text{Var}(x)$ that are induced by daily variation in news arrival A_t (through A_t 's impact on $\text{Var}(e)$ and $\text{Var}(\xi)$) are sufficient to identify all the parameters of the model.¹⁷

¹⁷ The covariance equation above also clarifies why positive values of α or ϕ could in principle account for the positive *unconditional* covariance between price changes and order flow found in the data. The model estimates allow us to determine whether this covariance arises because order flow has a positive price impact or because price changes feedback positively into order flow. We also estimated a version of the model where total order flow x is included in equation (5), as opposed to the non-feedback-driven

There are at least two reasons why contemporaneous returns might enter the daily order flow equation.¹⁸ The first is aggregation bias. As we show in the Appendix, when the intraday relationship between order flow and returns has an extended intertemporal structure, then the simple sums that daily data represent cannot capture the relationship fully. In this case, time-aggregated order flow x_t is not quite the right measure to include as a regressor in equation (1). As a result, joint estimation of equations (5) through (8) will produce a non-zero ϕ . The second reason why contemporaneous returns might enter the daily order flow equation is intraday feedback trading.¹⁹ (In daily data, intraday feedback trading would appear contemporaneous.) If intraday returns do not follow a martingale then individuals might rationally choose a feedback strategy. For example, if intraday returns exhibit reversals (and they do, as we show), then negative feedback trading would be the rational response.

Though our specification in equations (5)-(8) address reverse causality (from returns to order flow), there is another causality issue we have not yet addressed: Can positive macro news cause an increase in both order flow and price (with no causality from order flow to price)? This is essentially the same question we addressed above when examining whether the common knowledge effect from macro news ξ_t should be orthogonal to our model's order flow innovations. Recall from that discussion that positive news that systematically increases signed order flow would be inconsistent with rational expectations. The reason is because—under rational expectations—public information is impounded in price instantaneously. At the new price, which fully impounds the public information, there is no longer

flow x^* . Estimates from this alternative specification are not substantially different from those we report in Table 2.

¹⁸ At the daily frequency, order flow and lagged returns in our data are uncorrelated, so we do not lose generality by including only contemporaneous returns in the order flow equation.

¹⁹ That feedback trading is even feasible on average, given that marketmakers are on both sides of every transaction in our data, is not obvious. Why wouldn't marketmakers anticipate these inter-marketmaker feedback trades and attempt to undo the anticipatable effects on their positions from the incoming orders? They would. But in equilibrium they would be unable to do so by trading with other marketmakers: somebody has to take the other side of these trades. For inter-marketmaker trading to be positive feedback on average, the trades of non-marketmakers would have to be negative feedback on average. This is possible, but not very satisfying intellectually. Osler (2002) reports evidence on the types of orders non-marketmakers actually submit. For evidence on the currency transaction flows of institutional investors, a specific class of non-marketmaker, see Froot and Ramadorai (2002).

motivation for dollar buying relative to dollar selling.²⁰ The change in price level may induce trading at the new price (i.e., unsigned volume). But on average good news for the dollar will not produce positive order flow in the aggregate at the new price (a relative increase in transactions initiated by dollar buyers).

The theory on which we base our empirical models is subtler. Part of the public news maps directly into price, as in the traditional story (the shock ξ_t). But if to some degree market participants are drawing different inferences from common macro data, then price-setting marketmakers need to clear the market by determining the inferences of others (which they cannot know a priori). How do they learn them? The answer from microstructure theory is that they learn from the sequence of submitted orders over time. In this case, price adjusts instantaneously to the marketmaker's rational expectation of the market's interpretation, and then goes through a period of gradual adjustment caused by the sequence of transacted orders.

Daily Results

We turn to the empirical model in equations (5) through (8) for sorting out the direct versus indirect (via order flow) effects of news. The full empirical model embeds six parameters, α , ω , σ , ϕ , s_κ , and s_η . The parameter α captures the price impact of order flow. For the central question of this paper—how is macro news transmitted to exchange rates—the key parameters are ω and σ . ω governs the direct effect of news on price: the greater the number of news arrivals A_t the higher the variance of this component of daily returns. σ governs the direct effect of news on order flow: the greater the number of news arrivals A_t the higher the variance of this component of daily order flow.²¹ The variance s_κ reflects the component of daily

²⁰ One way to produce this non-causal correlation between order flow and price changes is if marketmakers are able to execute dollar purchases against other marketmakers before price adjusts to the public information. But this would require irrationality on the part of the quoting marketmaker: by not instantaneously impounding the public information in price he is permitting trades at “stale” prices, clearly a losing proposition.

²¹ Though the derivation for model equation (1) does not call for order flow lags, whether they are relevant empirically is a legitimate robustness concern. The daily analysis in Evans and Lyons (2002a) shows, however, that order flow lags are insignificant when included in this specification. Indeed, our

returns not explained by order flow or news. The variance s_η reflects the component of daily order flow not explained by news.

Table 2 reports GMM parameter estimates together with standard errors calculated from the asymptotic covariance matrix (allowing for heteroskedasticity).²² The table reports estimates for specifications using the flow of all news arrivals lumped together (labeled All News) and also with the U.S. and German news introduced separately. In both specifications the estimate of the price-impact parameter α is positive, as the theory predicts, and statistically significant. (Its size corresponds to a price impact of roughly 50 basis points per \$1 billion in order flow.) The parameter ϕ , which captures the contemporaneous effect of returns on order flow, is negative and significant. Recall that we offered two reasons why a negative relationship might exist here: temporal aggregation bias and negative feedback trading (on average) by marketmakers. We address these hypotheses in more detail in our intraday analysis. (At this stage, let us simply add that it is comforting that these results regarding news effects are robust to including a role for ϕ . The overall news results are essentially unchanged when we restrict ϕ equal to zero.) With all news arrivals lumped together, both of the key parameters ω and σ are significant and correctly signed (positive), implying that both direct and indirect effects of news on price are present. When U.S. and German news events are introduced separately (last column), both of the ω estimates are significant and correctly signed (positive). Also, the effect of German news on order flow volatility (σ_g) is positive and significant. The effect of U.S. news is of similar size but is not significant at the five percent level. (Recall that the median number of daily German news arrivals is nearly three times that for U.S. news arrivals.) In both specifications, the two (unconstrained) variance parameters s_κ and s_η are quite significant. The next rows of the table report Wald statistics for various parameter restrictions. The first two of these rows show that in the unrestricted model (i.e., with U.S. and German news

order flow measure shows no persistence in daily data, nor do daily returns, so we are not omitting any variables that a non-structural VAR approach would identify as significant.

²² Though not reported, we estimated the model allowing the price-impact parameter α to vary linearly with the number of announcements: the coefficient on the incremental announcement effect was positive but not significant. We also tested whether announcement effects on the variance of order flow might be non-linear (by modeling the variance of e_t as $\sigma A_t'$), but found no evidence of this either.

separated), the null that the direct and indirect news channels are insignificant are strongly rejected. From the third row, the restriction that both of the ω and σ coefficient pairs are equal cannot be rejected, which provides support for the restricted specification in the All News column.

Important summary statistics are provided in the final rows of Table 2. The three rows labeled with $R_{\Delta p}^2$ present the fraction of total variance in price attributable on average to the direct and indirect effects of news. In both cases (“All News Together” and “US/German News Separate”) the indirect channel through order flow is roughly twice the size of the direct effect: about 20 percent versus about 10 percent, respectively.²³ The final row presents the P-value for the test that the direct effect is larger than the indirect effect. In the restricted (All News Together) model, the hypothesis that the direct effect is larger is rejected at about the 1 percent level. (Recall that the restrictions on that model are not rejected.)

To summarize, we do indeed find that order flow helps transmit macro news to exchange rates. Roughly 20 percent of the daily movements in spot rates can be linked to macroeconomic news through the indirect order-flow channel. Through the direct (traditional) channel, macro news accounts for about 10 percent of daily price movements.²⁴ Given rejection of the test for whether $R_{\Delta p}^2(\text{direct}) > R_{\Delta p}^2(\text{indirect})$, we conclude that at least half of the effect of macro news on the exchange rate is transmitted via order flow. Finally, let us provide a big-picture accounting from these point estimates. The analysis in Evans and Lyons (2002a) shows that about 60 percent of daily DM/\$ price variation is due to order flow and about 40 percent is due to other factors. The results in Table 2 shed light on both of these pieces. They suggest that order flow’s 60 percent breaks into 20 percent that is induced by macro

²³ Given that much of the past work on news and exchange rates implicitly assumes that the indirect channel is not present, we also estimate our daily model with the (rejected) restriction that σ equals zero. We find that the $R_{\Delta p}^2$ from the direct channel rises above 10 percent, but only to 13.6 percent, i.e., it does not rise to anywhere near the 29 percent we find for both channels. This may help explain why the explanatory power of news in our model is relatively high (see too the analysis below of the subset of scheduled announcements).

²⁴ Noise in our measurement of news arrivals may bias down our estimates of the total percentage of price variance from news flow, but it is not obvious why this should bias our estimates of the *relative* importance of the direct and indirect channels. With respect to the possible biasing down of the total percentage, note that our results are positive in this respect (i.e., we find a large percentage relative to existing literature); eliminating the bias, if it exists, would only make our result more striking.

news and 40 percent that is not news induced. The 40 percent due to other factors in that earlier work breaks into 10 percent from the direct effect of macro news and 30 percent that remains unaccounted for.

Robustness

In this subsection we consider several potential robustness issues. First, we provide some further analysis of the feedback-trading hypothesis. Suppose one thought that observed order flow is the sum of exogenous order flow x_t^* , and a “feedback” component proportional to the change in price expected in the *future*:

$$x_t = x_t^* + \phi E[\Delta p_{t+1} | \Delta p_{t-1}]$$

In this case, measured order flow might respond to past price changes because those changes lead to revision in feedback traders’ forecasts of future price changes. One way to test whether this ϕ equals zero is to consider the instrumental variables (IV) regression:

$$x_t = \phi \Delta p_{t+1} + a_1 x_{t-1} + a_2 x_{t-2} + \dots + v_t$$

where Δp_{t-1} can be used to instrument for Δp_{t+1} . When estimated, we find that the IV ϕ estimates are insignificantly different from zero. Though not obvious, this is in fact what one would expect because in these data future price changes are unpredictable based on past price changes alone (see the autocorrelations in Table 1). There is simply no evidence here to suggest that past price changes should generate current order flow because they forecast future returns.

Our second robustness issue concerns the following question: What role does our broad definition of news play in accounting for the relatively large share of volatility due to public news arrival (33 percent)? To address this we turn to the subset of our news data that represents scheduled macroeconomic announcements, the subset for which we have ex-ante survey expectations for measuring announcement innovations. We focus on four U.S. announcements in particular, all of which are consistently among the most important in past event studies: Non Farm Payroll,

Durable Goods, Trade Balance, and Unemployment Claims.²⁵ We create a set of standardized forecast errors using the sample standard deviations of the measured innovations. Then, for each day in the sample we construct an index A_t equal to the sum of the absolute standardized forecast errors. Then we re-estimate our model in Table 2 using this alternative measure of macro news flow. The results from this version of the model are in Appendix Table 1.

The results are similar in some ways to those in Table 2, but quite dissimilar in other ways. They are similar in that the parameter ϕ is still negative and still significant. The key parameters ω and σ are also still positive and significant (at the 4 and 6 percent levels, respectively). The big difference is in the bottom line: the variance ratios are 0.014 for the direct channel and 0.028 for the indirect channel (All News Together case), implying that these scheduled announcements account for only about 4 percent of total exchange rate variance. Thus, using this more limited measure of news accounts for far less of the daily price variance than the full set of Reuter’s Money Market Headline News. This is consistent with past work, which shows that scheduled announcements account for less than 10 percent of daily price variance. Thus, it does appear that our broad definition of macro news can account for our finding that a large share of volatility is due to public news arrival.

Appendix Table 1 also speaks to a related robustness check: might our definition of macro news account for the indirect channel (via order flow) being twice as important for exchange rate variance as the direct channel (i.e., roughly 20 of 30 percent)? Note that the ratio of indirect to direct effects in Appendix Table 1 remains about the same as those in Table 2, namely 2 to 1, suggesting that the relative importance of the indirect channel is not due to our definition of news.

As a final type of robustness check we turn to regression analysis to examine the sources and form of heteroskedasticity. We consider (i) the extent to which news flow can account for the conditional variance of returns and order flow and (ii) whether the linear specifications of the variance functions $\sigma(A_t)$ and $\omega(A_t)$ in our

²⁵ We select these announcements to give the “scheduled announcement” model a greater chance of success. In fact, results from including all the scheduled U.S. announcements for which survey expectations are available are even poorer than the results we report, presumably because most of these other announcements are insignificant for exchange rates (as shown in earlier work). For German announcements, there is little empirical basis for selecting the more significant announcement types so we include all those announcements for which we have survey expectations.

baseline model are warranted. As described in the Appendix, we test for news-related heteroskedasticity by regressing the daily conditional variance of returns on a constant and the daily news flow A_t . We perform the same regression for the daily conditional variance of order flow. We find a significant direct effect of news on the conditional variance of prices (see Appendix Table 2). We also find a significant direct effect of news on the conditional variance of order flow. The first of these two results is consistent with existing literature. The second is new, though, and quite important to the economics underlying our model. (Both results are consistent with the visual evidence in the scatter plots in Figure 3.) Regarding potential non-linearity in the news flow effects, we find no evidence that squared news-flow terms should be included in the regression.

To summarize the robustness results above, we find that our broad definition of macro news is important for explaining a large fraction of exchange rate variation. Our news definition is also important for understanding why the model explains so much exchange rate variation relative to existing literature. In contrast, our news definition does not appear to be important for the result that the indirect macro news channel is at least as important as the direct channel. Our regression analysis finds a strong effect of news flow on order flow volatility: the order flow process is not orthogonal to public news, as simple rational expectations models would predict. Finally, our regression results support the linear specifications we adopt in our GMM model and corroborate that both direct and indirect price effects from news are present.

3. Intraday Analysis: Direct vs Indirect Effects from News

Our second approach complements the daily analysis above by examining intraday data, specifically five-minute data. At this ultra-high frequency, we have sufficient power to estimate both returns and order flow as state-dependent processes. These processes allow the arrival of news to affect price dynamics quite generally. Also, our intraday analysis allows us to estimate the percentage of price variation at different horizons that comes from order flow shocks, and how that percentage changes with news arrival. This is important for answering whether order flow plays a larger or smaller role in exchange rate determination when news

arrives. Finally, unlike our daily analysis, which permitted only unconditional inferences, our intraday analysis allows us to draw conclusions about news, order flow, and price conditional on trading intensity and price-impact horizon.

Our intraday model, designed for use in five-minute data, is based on the same underlying economic environment as our daily model. In particular, it extends the empirical model in Evans (2002). This model takes the form:

$$\Delta p_i = B(L)v_i + \varepsilon_i \quad (9)$$

$$x_i = \tilde{C}(L)v_i \quad (10)$$

where now we use the subscript i to denote the five-minute frequency. Here, v_i represents dispersed information shocks, i.e., bits of information that are first manifested in order flow and then subsequently impounded in price. The dynamic responses of prices and order flow to these dispersed information shocks are determined by the lag polynomials $B(L)$ and $\tilde{C}(L)$. The other term in the price equation corresponds to common-knowledge news ε_i . The v_i and ε_i shocks are mutually independent and serially uncorrelated.

There are two key identifying assumptions: rational expectations (as before) and imperfect substitutability across different-currency assets. To understand the first, note that we require orthogonality between the common-knowledge news shock ε_i and the dispersed information shock v_i . This orthogonality derives from the rational expectations restriction that common-knowledge news shocks are impounded fully and instantaneously in price (see the discussion of causality within the daily model). The second identifying assumption maintains that all order flow x_i represents dispersed information. Imperfect substitutability together with rational expectations is sufficient to ensure this: with imperfect substitutability (i.e., aggregate foreign exchange demand that is imperfectly elastic), all order flow conveys information that is price relevant since all order flow requires price adjustment to achieve market clearing.²⁶ (The rational expectations assumption

²⁶ In the extreme of perfect substitutability, in contrast, order flow *can* include components that are price irrelevant. It is perhaps best to think of the dispersed information shock v_i in equation (10) as a

ensures this information is dispersed rather than public.) Based on this specification in equation (10) we shall address whether our intraday estimates are consistent with an aggregation-bias explanation for the negative estimates of ϕ (reported in Table 2).

In ultra-high frequency data, such as the five-minute data we examine in this section, it is important to account for subtleties in the timing of dispersed information shocks relative to the timing of order flow. In particular, the first event in the process of foreign exchange trading is a desire to transact on the part of the non-marketmaker customer. This is not the order flow we measure, however. The trades between marketmakers measured by x_i are temporally downstream from the initial customer order. If, for example, some information conveyed by customer orders is impounded in price when customer orders are executed, then price is adjusting *before* trades between marketmakers are induced, despite the fact that order flow—the customer order—is driving the price change (see Evans 2001, pages 48-49). It is important to account for these initial stages of price adjustment. Figure 4 provides additional perspective. The bottom line is that even when causality is running wholly from order flow to price, the temporal link between our order flow measure x_i and price changes will involve both leads and lags.

To develop an estimable model, we proceed in two steps. First, we assume that the lag polynomial in the order flow equation, $\tilde{C}(L)$, can be written as $L^m C(L)$ where $C(L)^{-1}$ exists (an autoregressive process). m denotes the number of (five-minute) periods between the first effects on price from a typical customer order and its subsequent effect on order flow between marketmakers, x_i . With these assumptions we can rewrite equation (9) as

$$\Delta p_i = D(L)x_i + \varepsilon_i, \quad (11)$$

where $D(L) = B(L)L^{-m}C(L)^{-1}$. The polynomial $D(L)$ may take many forms depending on the dynamic responses of price and order flow to dispersed information

composite shock that includes any information relating to narrowly defined fundamentals as well as information relating to portfolio balance effects on price.

shocks. Notice though that if $m > 0$ (consistent with the timing illustrated in Figure 4) then $D(L)$ should contain both leads and lags (because the polynomials $B(L)$ and $C(L)^{-1}$ contain only non-negative powers of L). Based on the many tests presented in Evans (2002), we follow that paper in using a six-term polynomial: $D(L) = d_1L^{-4} + d_2L^{-3} + \dots + d_6L$ as our basic specification. Similarly, we find that the dynamics of order flow are well characterized if $C(L)^{-1}$ is an AR(10). In both cases, coefficient estimates on further leads/lags for price changes, and lags for order flow, are not statistically significant when more general specifications are estimated (see Evans 2002 for further diagnostics supporting these specifications). The specification for $D(L)$ implies that a customer order can begin to have an impact on price approximately 20 minutes before it affects order flow between marketmakers (i.e., $m = 4$).

The second step is to test whether the lag polynomials and the error variances are dependent on macro news having arrived. Evans (2002) reports strong evidence of dependence on trading intensity, n_i , measured as the number of transactions in observation window i . (Including trading intensity as a state variable is important for accommodating the pronounced time dependence in volatility documented by Andersen and Bollerslev 1998.) Here we extend the analysis by including macro news arrival as an additional state variable. Specifically, we construct dummy variables A_i^* and A_i that take the value of one if there was a news arrival (either U.S. or German) during the previous 15 or 5 minutes respectively.²⁷ The results in Appendix Table 3 shows that there is strong evidence for state-dependence in the return equation polynomial $D(L)$, with respect to both A_i^* and n_i , but not the order flow polynomial $C(L)^{-1}$. There is also strong evidence of heteroskedasticity related to A_i and n_i in the error variances.

²⁷ As noted, we use a dummy variable approach here because there are very few instances where there is more than one news arrival during a single 5-minute observation window. Note too that for our intraday analysis there is no scope for using the MMS-measured expected values for scheduled announcements (to measure announcement innovations) because even for monthly announcements there are only four forecast errors in our sample for each announcement type.

These results point to the need to incorporate state-dependence into our intraday model. To this end, we consider an extension of (9) where $D(L)$ is replaced by $D(L, n_i, A_i^*)$, a state-dependent 6'th order polynomial:

$$D(L, n, A^*) = d_1(n, A^*)L^{-4} + d_2(n, A^*)L^{-3} + \dots + d_6(n, A^*)L,$$

with state-dependent coefficients $d_j(n, A^*)$. Thus, $d_6(n, A^* = 1)$ is the coefficient on lagged order flow when trade intensity equals n and news arrived in the past 15 minutes ($A^* = 1$). We also allow for heteroskedasticity in the error variances, $Var(\varepsilon_i) = \Sigma_\varepsilon(n_i, A_i)$ and $Var(v_i) = \Sigma_v(n_i, A_i)$. State-dependence in the coefficients and variances is modeled as:

$$d_j(n, A^*) = d_j(0, A^*)\exp(-n/\gamma) + d_j(\infty, A^*)(1 - \exp(-n/\gamma))$$

$$\Sigma_j(n, A) = \Sigma_j(0, A)\exp(-n/\gamma) + \Sigma_j(\infty, A)(1 - \exp(-n/\gamma))$$

where $d_j(0, 0)$, $d_j(\infty, 0)$, $\Sigma_j(0, 0)$, and $\Sigma_j(\infty, 0)$ are the parameters to be estimated for observations without a news arrival, and $d_j(0, 1)$, $d_j(\infty, 1)$, $\Sigma_j(0, 1)$, and $\Sigma_j(\infty, 1)$ when there is a news arrival. Although these functional forms are somewhat specialized with respect to variations in trading intensity, they do not restrict how the flow of news affects price and order flow dynamics. Nor do they appear unduly restrictive when we subject our model to specification tests below.

Intraday Results

Table 3 presents GMM estimates of the intraday model. One measure of the importance of state-dependence in the price change dynamics is the sum of the order flow coefficients $D(1, n, A^*)$. This measures the long run impact of order flow on the price level. The estimated $D(1, n, A^*)$ varies considerably with both trading intensity and macro news flow. This is consistent with the findings of our state-dependence tests. It also accords well with the non-parametric evidence on state-dependence in hourly price change data reported in Evans and Lyons (2002b). The

variation in the estimates of $D(1, n, A^*)$ also underlines how the analysis of intraday data brings greater resolution to the study of return and order flow dynamics.

Before addressing the implications of our estimates, we examine a number of diagnostic tests (see Appendix Table 4). First, we examine the null hypothesis of no state-dependence in $D(1, n, A^*)$. Consistent with our earlier results, this hypothesis is strongly rejected with respect to trade intensity and macro news flow. Second, we consider LM-type tests for mis-specification in the estimated $d_j(n, A^*)$ and $\Sigma_j(n, A)$ functions. None of these tests are statistically significant suggesting that the estimated model has captured well the state-dependence related to trading intensity and news flow. Third, we examine whether the model captures all the heteroskedasticity in the data. Here we find evidence of some residual ARCH effects. So, while the model does appear to capture the state-dependent dynamics of prices and order flow with respect to trading intensity and news flow, it does not account for all the sources of heteroskedasticity. As a final diagnostic, we examine whether the estimated intraday model can account for the negative estimates of ϕ found in the daily model. Under the null of our intraday model, time-aggregated (i.e., daily) order flow x_t is only approximately equal to the correct measure of order flow that should be present in equation (1). Moreover, as the appendix shows, estimates of the intraday model imply that the approximation error is negatively correlated with daily price changes. In other words, the negative estimates of ϕ can arise through time-aggregation of the price and order flow dynamics that are described by our intraday model.

We now turn to the point of this section: How is the information in macro news transmitted to prices? To answer this question we use the GMM estimates to compute the fraction of price variance attributable to dispersed information across different market states. Specifically, our estimates allow us to write price changes as:

$$\Delta p_i = B(L, n_i, A_i^*)v_i + \varepsilon_i$$

where $B(L, n, A^*) = D(L, n, A^*)L^4C(L)$. Using this equation, we can compute the fraction of the variance in a k-period price change due to v_i shocks for a given level

of trade intensity and news flow. Table 4 shows that this fraction increases with trade intensity and horizon (as in Evans 2002). More importantly, however, there is a marked difference in the contribution depending on whether news arrives. For all trade intensities and horizons, dispersed information contributes more to the variance of price changes when news has arrived. Indeed, comparison of the first and second panels shows that order flow accounts for two to three times as much price variance when news has arrived. Thus, order flow is *more* important in exchange rate determination when news arrives, not less. The p-values for the hypothesis that there is no increase in the ratio, shown in the lower panel, are all small, particularly at the 5-minute horizon and in the limit as the horizon approaches infinity.

Robustness

In this subsection we consider several potential robustness issues at the intraday frequency. The first concerns the possibility of feedback trading. The last section of the appendix addresses whether in this empirical return environment feedback trading could possibly produce profits on average. There we consider feedback trading rules of a rather general form (equation A5) and examine profit outcomes from Monte Carlo simulations. Based on typical trade sizes in this market and reasonable levels of trading intensity, we find that strategies within the broad class we consider produce trading losses of about \$1 million per month on average, with probabilities of monthly losses of about 60 percent. Moreover, these estimates are conservative in that they take no account of transaction costs (neither in terms of price impact nor spreads).

That feedback trading would not be profitable in this environment does not (for many people) rule out its presence. For these people we construct a lower bound on price variance due to order flow shocks by including only those price effects from flow that occur after flow. (We do this despite the message from Figure 4 that some order flow between dealers *should* lag price adjustment, even when causality is running wholly from order flow—end-user order flow—to price.) These lower-bound results are presented in Appendix Table 5 (following the same format as Table 4). As expected, by this measure order flow contributes less to the variance of prices

unconditionally. But there is still strong evidence that order flow contributes more to the variance of prices in the presence of macro news. The essential finding in Table 4 is therefore robust to extreme views on the prevalence of feedback trading.

For some, another less satisfying feature of our analysis is that it is a significant departure from the existing event-study-based literature. To provide more of a bridge, and also to assuage concerns that order flow's behavior may bear no relation to event study logic, we provide an additional form of evidence. This evidence is informal, however, as it is based on that small subset of our sample for which we can measure ex-ante expectations, namely, events defined by scheduled macro announcements. Figure 5 presents the average price and order flow responses to two such announcements, shown in past work to be particularly relevant for exchange rates: the Trade Balance and Unemployment Claims (both U.S.). The figure shows the average cumulative response of order flow (solid line) and log prices ($\times 10^4$, dashed line) around a +1 standard deviation surprise in these announcements. (The horizontal scale is the number of 5-minute periods since the period containing the announcement.) Though informal, the plot makes the point that order flow moves dramatically, and in the same direction as prices. In the case of the trade balance, for example, improvement is on average good for the dollar, as the price here is measured as the DM price of a dollar. With respect to response time, remember that these are scheduled announcements, so participants can plan their responses (conditional on realizations) in advance, which lowers response lags considerably relative to unscheduled news. In any event, recognize that these are average responses here. In response to news, order flow and price are also (jointly) more volatile, which adds a lot to the first-moment story told by the plot.

4. Conclusions

This paper extends the literature on exchange rates and public news in three main ways. First, our analysis is the first to address the presence and importance of an indirect channel through which public news affects exchange rates. Second, methodologically we depart from existing literature on exchange rates and public news by using heteroskedasticity-based identification, a la Rigobon and Sack (2002), rather than the event-study approach. Third, our methodology accommodates a much wider set of news events than the scheduled macro announcements that are

the focus of event studies.

To address the indirect channel, we test whether macroeconomic news is transmitted to exchange rates via induced transactions and if so, what share occurs via transactions versus the traditional direct channel. Our model distinguishes three sources of exchange-rate variation. The first source mirrors traditional models—public news that is impounded in price immediately and directly (i.e., with no role for order flow). The second source is an indirect effect of public news that operates via induced order flow. The third source of exchange rate variation is caused by order flow that is unrelated to public news arrival. No previous work has disentangled these three sources empirically.

Using DM/\$ data from 1996, we find that all three sources of price variation are significant. Our point estimates at the daily frequency imply that about two thirds of the price effect from macro news is transmitted via order flow, with the remaining one third being the direct impounding of news in price. Unconditionally, the total effect accounts for about 30 percent of long-horizon exchange rate variance. We reject the null that the direct channel is the more important, implying that the indirect channel accounts for at least half of the effect of macro news on the DM/\$ rate.²⁸ These daily results are consistent with those at the five-minute frequency. In particular, our five-minute results in Table 4 show that when news arrives there is a 100-200 percent increase in the importance of order flow in price determination. This is inconsistent with the “macroeconomic” prior that order flow should be less important when public news arrives. These intraday results apply consistently when conditioning on both trading intensity and price-impact horizon.

Why is news affecting exchange rates via order flow? Our preferred interpretation is that participants in this market are drawing different inferences from common macroeconomic data. This could be rational if, for example, agents observe the common data with different prior beliefs (perhaps based on different information sets). In any event, whether this is rational or not is beyond the scope of this paper. We would simply emphasize that the phenomenon is not limited to the FX market: think, for example, of differences in analyst forecasts of aggregate equity

²⁸ Whether this result generalizes to less actively traded currencies is beyond the scope of our analysis.

values, where these forecasts, too, should be based on common economy-wide data (Harris and Raviv 1993).

Why does this paper find that roughly 30 percent of total volatility comes from news effects whereas past work finds that news effects account for less than 5 percent? We offer three possible reasons. First, we consider a much wider set of macro news arrivals. This explanation is consistent with our finding that when we estimate our model on scheduled announcements only (the more limited set used by others, for which survey expectations are readily available) we find they account for only 4 percent of total volatility. Second, we adopt a heteroskedasticity approach that mitigates a source of measurement error, in that it does not rely on measuring ex-ante news expectations (in contrast to the event-study approach). Third, we use a less restrictive model in the sense that we allow for an indirect channel and we allow flow responses to dispersed information to persist over extended periods. Event studies impose window lengths that constrain response times, typically less than one hour in intraday analysis. (In practice, response times for some participants surely extend beyond an hour.) This constraint might appear rather loose: past work finds that volatility changes from announcements dissipate quickly (e.g., half-lives less than an hour). But volatility does not go to zero when shocks to volatility dissipate; it settles back to a (statistical) mean. Volatility being at its mean does not imply that last hour's news is irrelevant for current trading. Rather, participants' macro views evolve continuously and trades induced by those evolving views can hit the market over extended periods.²⁹ This idea links rather naturally to the analysis in Andersen et al. (2003): they find that the impact effect of announcements on price—the first moment effect—is over quite quickly, whereas the effect on volatility is not immediate, with the full “impact” effect on volatility occurring only after about an hour (their Table 3). Future analysis might address whether volatility decay rates that differ across announcements are an indicator of relative difficulty in interpretation.

²⁹ Carlson (2002), for example, finds that in response to a public macro announcement, liquidity in an electronic interbank trading system (Reuters 2000-2) remained significantly below normal (and below its ex ante state) for about 2 hours. Even if *average* effects from news are reflected in prices quickly, as in Andersen and Bollerslev (1998) and Cheung and Chinn (2001), this does not imply that total effects are reflected quickly. Average effects may correspond to the direct, or rational expectations channel, which one would expect to be reflected more quickly than indirect, order-flow mediated effects.

It would be heroic indeed for us to assert that causality between order flow and exchange rates is strictly one directional. It is, almost surely, two directional, at least under some circumstances.³⁰ That said, as a theoretical matter the portfolio shifts model underlying our analysis provides a clear formalization of one-directional causality, even in the face of public information. As an empirical matter, we allowed for the potential presence of reverse causality but found no evidence of the most commonly offered form, namely positive-feedback trading. Moreover, existing micro evidence of causality running from order flow to prices in stock, bond, and FX markets is so strong that there is no case for believing that causality runs wholly or even largely in reverse (for micro evidence in stocks and bonds see the survey by Madhavan 2000). Rather, we have presented many types of broadly complementary evidence that, together, make a strong case that indirect effects of news, via transactions, do indeed affect exchange rates. This evidence includes different data frequencies, different horizons, different identifying assumptions, etc. The story that emerges from the ensemble of evidence is clear, new, and substantial: by opening the order-flow channel, the ability of news to account for total exchange rate variation rises from the 1-5 percent found previously to roughly 30 percent, a sizable step in resolving the distressing puzzle of missing news effects.

³⁰ One possible channel of reverse causality is when exchange rate changes are large enough to induce distress at some institutions, causing them to sell into falling markets (and vice versa) for risk management purposes.

References

- Andersen, T., and T. Bollerslev (1998), Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies, *Journal of Finance*, 53: 219-265.
- Andersen, T., T. Bollerslev, F. Diebold, C. Vega (2003), Micro effects of macro announcements: Real-time price discovery in foreign exchange, *American Economic Review*, 93: 38-62.
- Andersen, T., T. Bollerslev, F. Diebold, and P. Labys (2001), The distribution of realized exchange rate volatility, *Journal of the American Statistical Association*, 96: 42-55.
- Balduzzi, P., E. Elton, and C. Green (2001), Economic news and bond prices, *Journal of Financial and Quantitative Analysis*, forthcoming.
- Bank for International Settlements—BIS (2002), Central bank survey of foreign exchange market activity in April 2001. Publication of the Monetary and Economics Department, BIS (at www.bis.org).
- Berry, T., and K. Howe (1994), Public information arrival, *Journal of Finance*, 49: 1331-1346.
- Carlson, J. (2002), One minute in the life of the DM/\$: Public information in an electronic market, typescript, Purdue University, August.
- Cheung, Y., and M. Chinn (2001), Currency traders and exchange rate dynamics: A survey of the U.S. market, *Journal of International Money and Finance*, 20: 439-471.
- Chinn, M. (1991), Some linear and non-linear thoughts on exchange rates, *Journal of International Money and Finance*, 10: 214-230.
- Cornell, B. (1982), Money supply announcements, interest rates, and foreign exchange, *Journal of International Money and Finance*, 1: 201-208.
- DeGennaro, R., and R. Shrieves (1997), Public information releases, private information arrival, and volatility in the foreign exchange market, *Journal of Empirical Finance*, 4: 295-315.
- Eddelbuttel, D., and T. McCurdy (1998), The impact of news on foreign exchange rates: Evidence from high frequency data, typescript, University of Toronto.
- Ederington, L., and J. Lee (1995), The short-run dynamics of price adjustment to new information, *Journal of Financial and Quantitative Analysis*, 30: 117-134.
- Engel, C., and J. Frankel (1984), Why interest rates react to money announcements: An answer from the foreign exchange market, *Journal of Monetary Economics*, 13: 31-39.
- Evans, M. (2001), FX trading and exchange rate dynamics, NBER Working Paper 8116, February.
- Evans, M. (2002), FX trading and exchange rate dynamics, *Journal of Finance*, 57: 2405-2448.
- Evans, M., and R. Lyons (2002a), Order flow and exchange rate dynamics, *Journal of Political Economy*, 110: 170-180.
- Evans, M., and R. Lyons (2002b), Time-varying liquidity in foreign exchange, *Journal of Monetary Economics*, 49: 1025-1051.

- Fleming, M., (2002), Price discovery in the U.S. treasury market: The impact of order flow and liquidity on the yield curve, typescript, New York Federal Reserve Bank, September.
- Fleming, M., and E. Remolona (1999), Price formation and liquidity in the U.S. treasury market, *Journal of Finance*, 54: 1901-1915.
- Flood, R., and A. Rose (1995), Fixing exchange rates: A virtual quest for fundamentals, *Journal of Monetary Economics* 36, 3-37.
- Frankel, J., and K. Froot (1990), Chartists, fundamentalists, and trading in the foreign exchange market, *American Economic Review*, 80: 181-185.
- Frankel, J., and A. Rose (1995), Empirical research on nominal exchange rates, in G. Grossman and K. Rogoff (eds.), *Handbook of International Economics*, Elsevier Science: Amsterdam, 1689-1729.
- French, K., and R. Roll (1986), Stock return variance: The arrival of information and the reaction of traders, *Journal of Financial Economics*, 17: 5-26.
- Froot, K., and T. Ramadorai (2002), Currency returns, institutional investor flows, and exchange rate fundamentals, NBER Working Paper 9101, August.
- Glesjer, H. (1969), A new test for heteroskedasticity, *Journal of the American Statistical Association*, 64: 316-323.
- Glosten, L., and P. Milgrom (1985), Bid, ask, and transaction prices in a specialist market with heterogeneously informed agents, *Journal of Financial Economics*, 14: 71-100.
- Goodhart, C., S. Hall, S. Henry, and B. Pesaran (1993), News effects in a high-frequency model of the sterling-dollar exchange rate, *Journal of Applied Econometrics*, 8: 1-13.
- Green, C., (2002), Economic news and the impact of trading on bond prices, typescript, Emory University, May, *Journal of Finance*, forthcoming.
- Greene, William H. (1997), *Econometric Analysis*, Prentice Hall.
- Grossman, S., and J. Stiglitz (1980), On the impossibility of informationally efficient markets, *American Economic Review*, 70: 393-408.
- Hakkio C., and D. Pearce (1985), The reaction of exchange rates to economic news, *Economic Inquiry*, 23: 621-635.
- Hardouvelis, G. (1988), Economic news, exchange rates, and interest rates, *Journal of International Money and Finance*, 7: 23-25.
- Harris, M., and A. Raviv (1993), Differences of opinion make a horse race, *Review of Financial Studies*, 6: 473-506.
- Hasbrouck, J. (1991), Measuring the information content of stock trades, *Journal of Finance*, 46: 179-207.
- Hasbrouck, J., and D. Seppi (2001), Common factors in prices, order flows, and liquidity, *Journal of Financial Economics*, 59: 383-411.
- Isard, P. (1995), *Exchange Rate Economics*, Cambridge University Press: Cambridge, UK.
- Ito, T., and V. Roley (1987), News from the U.S. and Japan: Which moves the yen/dollar exchange rate? *Journal of Monetary Economics*, 19: 255-277.
- Jones, C., G. Kaul, and M. Lipson (1994), Transactions, volume, and volatility, *Review of Financial Studies*, 7: 631-651.
- Jones, C., and M. Lipson (1999), Execution costs of institutional equity orders, *Journal of Financial Intermediation*, 8: 123-140.
- Kandel, E., and N. Pearson (1995), Differential interpretation of public signals and trade in speculative markets, *Journal of Political Economy*, 103: 831-872.

- Killeen, W., R. Lyons, and M. Moore (2001), Fixed versus flexible: Lessons from EMS order flow, NBER Working Paper 8491, September.
- Klein, M. (1991), Managing the dollar: Has the Plaza Agreement mattered? *Journal of Money, Credit, and Banking*, 23: 742-751.
- Kyle, A. (1985), Continuous auctions and insider trading, *Econometrica*, 53:1315-35.
- Lyons, R. (2001), *The Microstructure Approach to Exchange Rates*, MIT Press, (chapters at faculty.haas.berkeley.edu/lyons).
- Madhavan, A. (2000), Market microstructure: A survey, *Journal of Financial Markets*, 3:205-258.
- Madhavan, A., M. Richardson, and M. Roomans (1997), Why do security prices change? A transaction-level analysis of NYSE stocks, *Review of Financial Studies*, 10: 1035-1064.
- Mark, N. (1995), Exchange rates and fundamentals: Evidence on long-horizon predictability, *American Economic Review*, 85: 201-218.
- Meese, R., and K. Rogoff (1983a), Empirical exchange rate models of the seventies, *Journal of International Economics*, 14: 3-24.
- Meese, R., and K. Rogoff (1983b), The out-of-sample failure of empirical exchange rate models, in J. Frenkel (ed.), *Exchange Rate and International Macroeconomics*, University of Chicago Press: Chicago.
- Melvin, M., and X. Yin (2000), Public information arrival, exchange rate volatility, and quote frequency, *Economic Journal*, 110: 644-661.
- Newey, W., and K. West (1987), A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica*, 55: 703-708.
- Osler, C. (2002), Currency orders and exchange-rate dynamics: An explanation for the predictive success of technical analysis, *Journal of Finance*, forthcoming.
- Payne, R. (1999), Informed trade in spot foreign exchange markets: An empirical investigation, typescript, London School of Economics, January, *Journal of International Economics*, forthcoming.
- Rigobon, R., and B. Sack, (2002), The impact of monetary policy on asset prices, NBER Working Paper 8794, February.
- Rime, D. (2000), Private or public information in foreign exchange markets? An empirical analysis, typescript, University of Oslo, March.

Table 1: Sample Statistics

A: Daily Data		News Arrivals					
		Δp_t	Δx_t	$ \Delta p_t $	$ \Delta x_t $	US: A_t^{us}	German: A_t^g
5%	-11.9	-308	0.1	5	0	2	2
25%	-3.8	-61	1.4	28	1	6	9
50%	0.3	8	3.7	83	2	8	11
75%	3.4	91	5.9	140	5	12	15
95%	6.9	186	11.9	319	7	18	21
Max	12.4	339	20.7	449	9	22	27
Min	-20.7	-449	0.0	0	0	0	0
Sdt. Dev.	5.9	136.4	3.8	97.9	2.2	5.1	6.0
Skewness	-0.8	-0.6	1.5	1.5	0.7	0.5	0.2
Kurtosis	3.8	4.5	6.1	4.9	2.6	3.0	2.6

B: 5-Minute Data		Δp_i	Δx_i	$ \Delta p_i $	$ \Delta x_i $	n_i
5%	-1.4	-9	0.0	0	2	
25%	-0.3	-2	0.1	1	6	
50%	0.0	0	0.3	3	12	
75%	0.3	3	0.7	5	21	
95%	1.3	9	1.7	12	44	
Max	5.0	69	7.9	72	212	
Min	-7.9	-72	0.0	0	2	
Sdt. Dev.	0.8	5.6	0.6	4.2	15.5	
Skewness	-0.2	0.1	2.2	3.2	3.4	
Kurtosis	7.4	12.6	10.8	22.9	23.2	

Autocorrelations (p-values)									
	Lag=1	2	3	4	5	6	12	18	24
Δp_t	-0.305 (0.00)	-0.010 (0.35)	-0.004 (0.76)	-0.003 (0.79)	-0.004 (0.68)	0.013 (0.23)	0.004 (0.69)	0.018 (0.06)	0.004 (0.64)
x_t	0.231 (0.00)	0.105 (0.00)	0.093 (0.00)	0.077 (0.00)	0.060 (0.00)	0.058 (0.00)	0.027 (0.00)	0.023 (0.01)	0.004 (0.65)

Notes: Δp_t is 1000 times the change in the last DM purchase price for dollars between 4:00 pm on day t and day $t-1$. Δx_t is the total interdealer order flow over the same time interval. Δp_i and x_i are the corresponding price changes and order flows during 5-minute interval i , and n_i is the total number of trades. A_t^{us} and A_t^g are respectively the number of macro news arrivals observed on the Reuters Money Market Headline News screen relating to the US and Germany between 4:00 pm on day t and day $t-1$. Autocorrelations are computed from five-minute data by GMM, and the p -values are calculated from Wald tests of the null hypothesis of a zero correlation (allowing for conditional heteroskedasticity).

Table 2: Daily Model Estimates

$\Delta p_t = \alpha x_t^* + \xi_t + v_t \quad \text{with} \quad \text{VAR}(\xi_t) = \omega A_t, \quad \text{VAR}(\kappa_t) = s_\kappa$ $x_t = x_t^* + x_t^+ \quad \text{with} \quad x_t^* = e_t + \eta_t, \quad \text{VAR}(e_t) = \sigma A_t, \quad \text{VAR}(\eta_t) = s_\eta, \quad x_t^+ = \phi \Delta p_t$				
	All News Together		US/German News Separate	
Parameters	Estimate	Std. Err.	Estimate	Std. Err.
α	0.031	(0.002)	0.031	(0.002)
ϕ	-2.051	(0.172)	-2.048	(0.172)
s_κ	81.375	(10.714)	78.729	(9.761)
s_η	3.795	(0.502)	3.774	(0.493)
ω	2.791	(0.690)		
σ	0.168	(0.034)		
ω_{us}			4.896	(1.781)
ω_g			2.342	(0.912)
σ_{us}			0.176	(0.090)
σ_g			0.166	(0.039)
<u>Wald Tests:</u>			<u>Statistic</u>	<u>P-value</u>
$\omega_{us} = \omega_g = 0$			68.081	(0.000)
$\sigma_{us} = \sigma_g = 0$			25.230	(0.000)
$\omega_{us} = \omega_g \ \& \ \sigma_{us} = \sigma_g$			2.069	(0.355)
<u>Variance Ratios</u>	<u>Ratio</u>	<u>Std. Err.</u>	<u>Ratio</u>	<u>Std. Err.</u>
$R_{\Delta p}^2 (direct)$	0.094	(0.037)	0.126	(0.036)
$R_{\Delta p}^2 (indirect)$	0.232	(0.062)	0.205	(0.054)
$R_{\Delta p}^2 (total)$	0.326	(0.069)	0.331	(0.068)
P-value: direct>indirect		(0.013)		(0.080)
<p>Notes: The table reports GMM parameter estimates and asymptotic standard errors (corrected for heteroskedasticity) in parentheses. The Wald statistics are for the null hypothesis listed; p-values are reported in parentheses. The standard errors reported for the variance ratios are computed from a Monte Carlo experiment with 5000 replications. $R_{\Delta p}^2 (direct)$ and $R_{\Delta p}^2 (indirect)$ are the fraction of the daily variance in prices attributed to news arrival via the direct and indirect channels. The p-value in the last line is for the null hypothesis that $R_{\Delta p}^2 (direct) > R_{\Delta p}^2 (indirect)$ and is calculated from the Monte Carlo experiment.</p>				


Table 3: Intraday Model Estimates (GMM)

$\Delta p_i = D(L, n_i, A_i^*)x_i + \varepsilon_i + \omega_i - \omega_{i-1}$ with $Var(\varepsilon_i) = \Sigma_\varepsilon(n_i, A_i)$, $Var(\omega_i) = \Sigma_\omega(n_i, A_i)$ $C(L)^{-1}x_i = v_i$ with $Var(v_i) = \Sigma_v(n_i, A_i)$						
Price Eq.: No News Arrival	x_{i+4}	x_{i+3}	x_{i+2}	x_{i+1}	x_i	x_{i-1}
Coefficients in $D(L, 0, A_i^* = 0)$	0.030 (0.024)	0.040 (0.051)	0.033 (0.197)	-0.061 (0.041)	-0.106 (0.022)	-0.046 (0.027)
Coefficients in $D(L, \infty, A_i^* = 0)$	0.118 (0.104)	0.214 (0.188)	0.489 (0.598)	0.705 (0.147)	-0.238 (0.062)	-0.027 (0.081)
	$\Sigma_\varepsilon(0,0)$	$\Sigma_\varepsilon(\infty,0)$	$\Sigma_\omega(0,0)$	$\Sigma_\omega(\infty,0)$	$D(1,0,0)$	$D(1,\infty,0)$
	0.000 (N/A)	0.010 (<0.001)	0.002 (<0.001)	0.000 (N/A)	-0.110 (0.030)	1.262 (0.105)
Price Eq.: News Arrival	x_{i+4}	x_{i+3}	x_{i+2}	x_{i+1}	x_i	x_{i-1}
Coefficients in $D(L, 0, A_i^* = 1)$	-0.086 (0.063)	-0.001 (0.059)	0.043 (0.044)	-0.198 (0.077)	0.081 (0.069)	0.001 (0.062)
Coefficients in $D(L, \infty, A_i^* = 1)$	0.475 (0.179)	0.095 (0.187)	0.321 (0.108)	0.791 (0.218)	-0.623 (0.162)	0.110 (0.149)
	$\Sigma_\varepsilon(0,1)$	$\Sigma_\varepsilon(\infty,1)$	$\Sigma_\omega(0,1)$	$\Sigma_\omega(\infty,1)$	$D(1,0,1)$	$D(1,\infty,1)$
	0.000 (N/A)	0.006 (0.002)	0.002 (<0.001)	0.000 (N/A)	-0.159 (0.084)	1.169 (0.227)
Order Flow Eq.	x_{i-1}	x_{i-2}	x_{i-3}	x_{i-4}	x_{i-5}	x_{i-6}
Coefficients in $C(L)^{-1}$	0.210 (0.014)	0.036 (0.013)	0.048 (0.012)	0.033 (0.012)	0.019 (0.011)	0.025 (0.011)
	x_{i-7}	x_{i-8}	x_{i-9}	x_{i-10}		
	0.015 (0.010)	0.017 (0.012)	-0.016 (0.010)	0.020 (0.008)		
	$\Sigma_v(0,0)$	$\Sigma_v(\infty,0)$	$\Sigma_v(0,0)$	$\Sigma_v(\infty,1)$		
	0.000 (N/A)	0.032 (0.002)	0.000 (N/A)	0.028 (0.037)		
<p>Notes: GMM estimates with asymptotic standard errors in parenthesis corrected for conditional heteroskedasticity and an MA(1) error term. (Coefficients and standard errors in the first two panels are multiplied by 100.) The state-dependent polynomial in the price equation is $D(L, n_i, A_i^*) = d_1(n_i, A_i^*)L^{-1} + d_2(n_i, A_i^*)L^{-2} + \dots + d_6(n_i, A_i^*)L^{-6}$ where A_i^* is a dummy variable equal to one if there was a macro news arrival during the previous 15-minutes (i.e. 3 observation windows). State-dependency in the coefficients and variances is modeled as $d_j(n_i, A_i^*) = d_j(0, A_i^*)\exp(-n_i/\gamma) + d_j(\infty, A_i^*)(1 - \exp(-n_i/\gamma))$, and $\Sigma_j(n_i, \lambda) = \Sigma_j(0, A_i^*)\exp(-n_i/\gamma) + \Sigma_j(\infty, A_i^*)(1 - \exp(-n_i/\gamma))$ where $\gamma = 100$, and A_i is a dummy variable equal to one if there was a news arrival during the previous 5-minutes.</p>						

Table 4: Percentage Price Variance From Order Flow Shocks

News Arrival					
	Horizon (minutes)				
	5	30	60	120	∞
<u>Trade intensity n</u>					
10	4.54 (1.50)	9.51 (5.80)	10.65 (7.33)	11.82 (9.05)	12.16 (9.59)
20	9.03 (2.72)	16.50 (8.85)	17.70 (10.52)	18.83 (12.18)	19.15 (12.65)
30	14.62 (4.32)	23.68 (11.57)	24.68 (13.27)	25.60 (14.88)	25.85 (15.33)
No News Arrival					
	Horizon (minutes)				
	5	30	60	120	∞
<u>Trade intensity n</u>					
10	1.46 (1.11)	2.66 (2.58)	2.57 (3.25)	2.46 (4.12)	2.42 (0.49)
20	3.65 (1.23)	7.28 (2.61)	7.75 (3.01)	8.23 (3.44)	8.37 (0.73)
30	6.55 (1.58)	13.39 (3.33)	14.49 (3.66)	15.52 (3.96)	15.80 (1.31)
P-values: No increase in variance ratio					
	Horizon (minutes)				
	5	30	60	120	∞
<u>Trade intensity n</u>					
10	0.02	0.04	0.06	0.07	0.00
20	0.00	0.02	0.04	0.06	0.01
30	0.00	0.04	0.07	0.10	0.02
Notes: Estimated from five-minute data (model in Table 3). The upper two panels report the percentage of price variance over various horizons due to order flow shocks, given trade intensity n , with and without concurrent news arrival. (From Table 1, the median number of trades n per five-minute interval is 12 and the 75 th percentile is 21.) The standard errors reported in parenthesis in the upper panels, and p-values shown in the lower panel are calculated from 5000 Monte Carlo draws from the estimated asymptotic distribution of the GMM estimates.					

Figure 1: Sample Distribution of News Arrivals




Notes: Distribution of all news arrivals by 30 minute interval starting at the interval centered at 6:30 BST and ending at the interval centered at 20:00 BST. Sample: May 1 to August 31, 1996. Data on news arrivals are from the Reuters Money Market Headline News screen.

Figure 2: Macro News Sample

Date	Time	News
05/01/96	13:05:22	MARCH U.S. LEADING INDICATORS SHOW ECONOMY EASING
05/01/96	14:00:50	U.S. MARCH CONSTRUCTION SPENDING ROSE 3.1 PCT
05/01/96	14:10:14	MARCH U.S. CONSTRUCTION SPENDING REBOUNDS STRONGLY
05/02/96	6:05:18	GERMAN MARCH IMPORT PRICES CLIMB 0.3 PCT M/M
05/02/96	8:33:10	BUNDESBANK DOES NOT PLAN NEWS CONFERENCE TODAY
05/02/96	9:48:20	GERMAN CALL MONEY FALLS BACK TO 3.30/40 PCT
05/02/96	10:50:08	BUNDESBANK LEAVES INTEREST RATES UNCHANGED
05/02/96	10:51:56	GERMAN MARCH INDUSTRIAL OUTPUT DATA DUE 1130 GMT
05/02/96	12:30:40	U.S. JOBLESS CLAIMS FELL IN LATEST WEEK
05/02/96	12:31:00	U.S. Q1 1996 REAL GDP ROSE 2.8 PCT
05/02/96	14:00:30	U.S. MARCH FACTORY ORDERS ROSE 1.5 PCT
05/02/96	15:01:54	U.S. Q1 GDP SURGE SEEN JUST A BLIP IN MODEST TREND
05/02/96	15:10:32	GERMAN EMPLOYER TO OPPOSE ANTI-WAGE DUMPING LAW
05/03/96	9:56:32	GERMAN CALL MONEY EASES SLIGHTLY AHEAD OF WEEKEND
05/03/96	12:30:38	U.S. APRIL NON-FARM PAYROLLS ROSE 2,000
05/03/96	12:31:00	U.S. MARCH PERSONAL INCOME ROSE 0.5 PCT
05/03/96	12:38:44	U.S. JOBLESS RATE LOWER BUT LABOR MARKET LOOKS WEAK
05/03/96	13:42:16	MARCH US INCOME DATA SHOW MODEST GROWTH, INFLATION
05/03/96	14:00:16	U.S. MARCH HOUSING COMPLETIONS ROSE 5.1 PCT

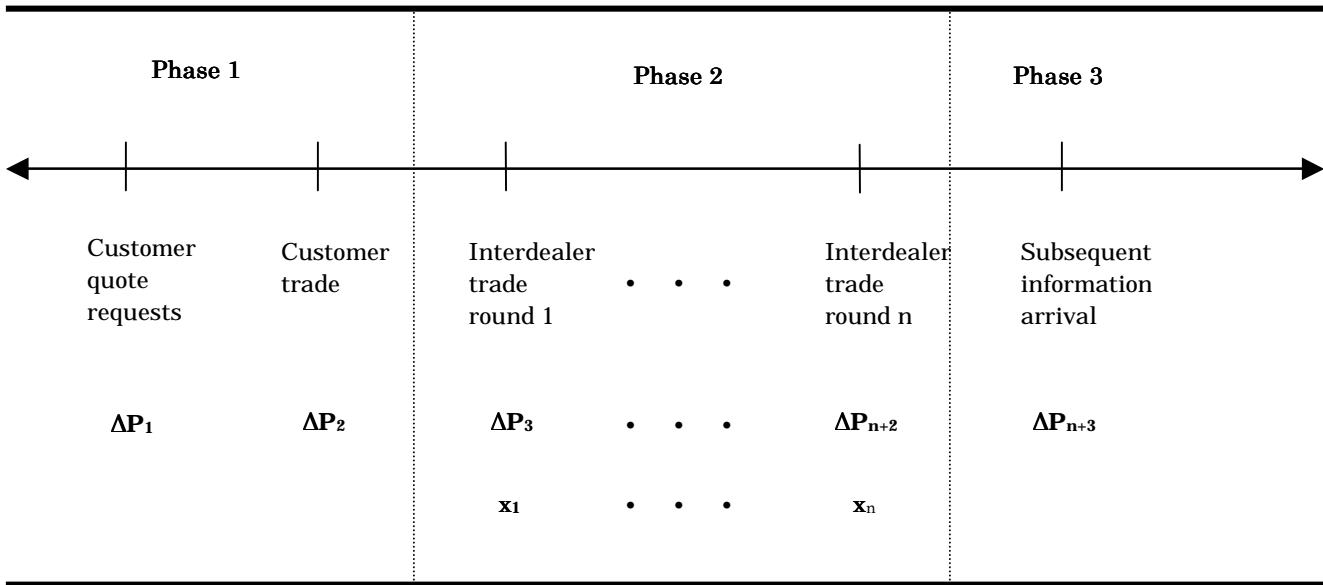
Notes: The figure shows the macro news arrivals on the first 3 days of our sample (May 1 to May 3, from the total sample from May 1 to August 31, 1996).

Figure 3: Daily Price and Order Flow Volatility



Notes: Daily data from May 1 to August 31, 1996. Data on news arrivals A_t are from the Reuters Money Market Headline News screen. Daily realized variances for Δp_t and x_t are integrated variance measures based on a five-minute sampling frequency (Andersen et al. 2001).

Figure 4: Time-line of Information Flow at Ultra-High Frequencies

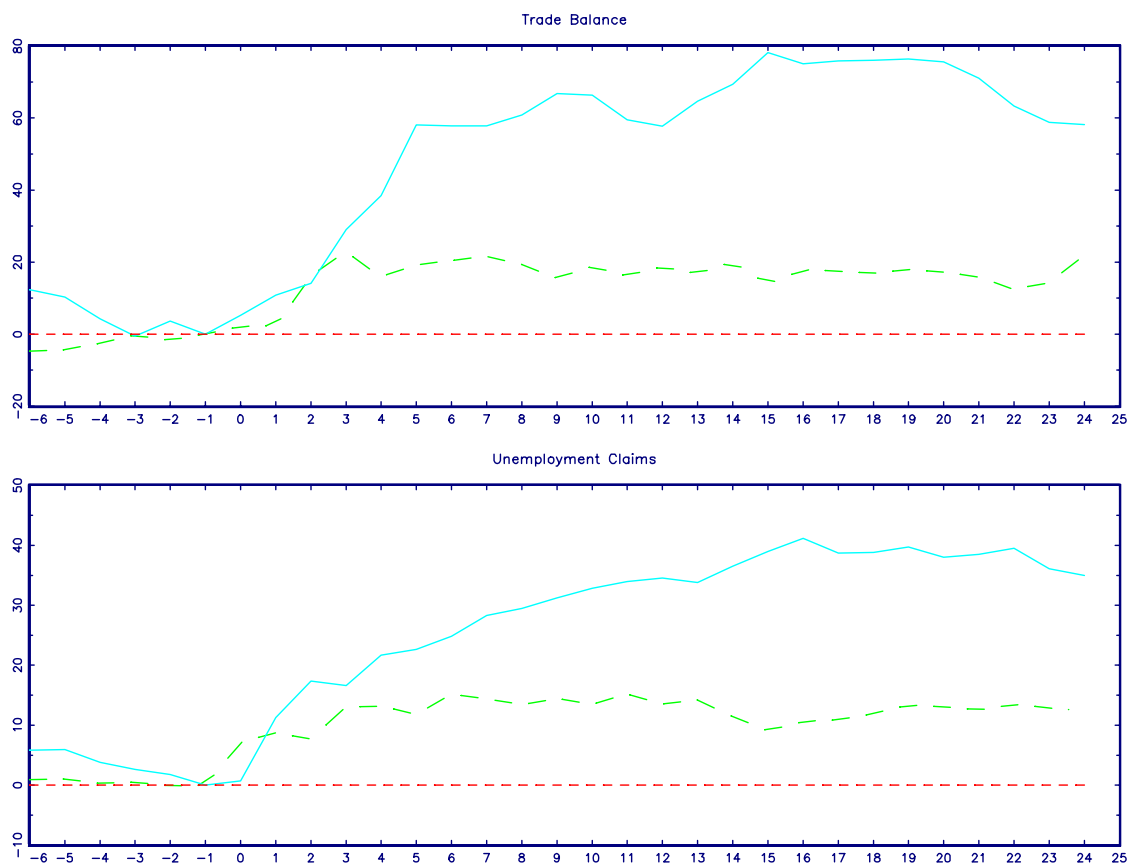


Notes: The response of price to order flow in ultra-high frequency data can be split into three phases, one where price adjustment precedes interdealer order flow, one where they are concurrent, and one where price adjustment follows interdealer order flow. (Recall that this paper’s empirical analysis uses order flow measured from interdealer trades.) In the first phase, price is responding to the orders received by dealers from non-dealer customers (ΔP_1 and ΔP_2). Price impact in this phase occurs before the induced interdealer trading (the latter involving both risk-sharing and speculative motives.) But all price impact does not occur in phase one because customer trades are not publicly observable. In the second phase, interdealer order flows x_i induced by the customer order—which are more transparent than the customer order—become the causal driver of additional price adjustment. The third phase includes price changes due to subsequent information arrival that is correlated with the initial customer trade (i.e., for which the initial customer trade provides a signal).

Rationales for the various stages of price adjustment come from multiple sources. The first price change shown above, ΔP_1 , is addressed by Jones and Lipson (1999): when customers call various banks for quotes, often the banks know something about the trade’s direction. So even banks that do not receive the order have learned something about it before it is executed, and some of this information becomes impounded in price before execution. The second of the phase-one price changes, ΔP_2 , is analyzed by Evans (2001, especially pages 48-49): when customer orders arrive, some adjustment in market price occurs immediately, i.e., before any interdealer trading; but, due to low transparency, not all. Turning to the phase-2 causal link between (rounds of) interdealer order flow and price changes ΔP_3 through ΔP_{n+2} , for detail see, e.g., Lyons (1997). The phase-3 link between price change and lagged order flow is standard within the large literature on information-based models of trading.

Figure 5

Price and Order Flow Responses to Two Scheduled Announcements



Notes: Figures show the cumulative (average) response of order flow (solid line) and log prices ($\times 10^4$, dashed line) around scheduled announcements of the US Trade Balance and Unemployment Claims. The responses are based on a +1 standard deviation announcement surprise. The horizontal scale shows the number of 5-minute periods since the period containing the announcement.

Appendix Table 1

Daily Model Estimates with Scheduled Announcements

$$\Delta p_t = \alpha x_t^* + \xi_t + v_t \quad \text{with} \quad \text{VAR}(\xi_t) = \omega A_t, \quad \text{VAR}(\kappa_t) = s_\kappa$$

$$x_t = x_t^* + x_t^+ \quad \text{with} \quad x_t^* = e_t + \eta_t, \quad \text{VAR}(e_t) = \sigma A_t, \quad \text{VAR}(\eta_t) = s_\eta, \quad x_t^+ = \phi \Delta p_t$$

Parameters	All News Together		US/German News Separate	
	Estimate	Std. Err.	Estimate	Std. Err.
α	0.031	(0.002)	0.031	(0.002)
ϕ	-2.036	(0.175)	-2.028	(0.174)
s_κ	108.725	(5.596)	108.582	(5.639)
s_η	5.506	(0.381)	5.477	(0.379)
ω	8.523	(5.019)		
σ	0.415	(0.267)		
ω_{us}			13.979	(11.636)
ω_g			7.815	(4.871)
σ_{us}			0.648	(0.664)
σ_g			0.379	(0.265)
<u>Variance Ratios</u>	<u>Ratio</u>	<u>Std. Err.</u>	<u>Ratio</u>	<u>Std. Err.</u>
$R_{\Delta p}^2(direct)$	0.014	(0.009)	0.018	(0.011)
$R_{\Delta p}^2(indirect)$	0.028	(0.020)	0.025	(0.018)
$R_{\Delta p}^2(total)$	0.042	(0.026)	0.044	(0.026)

Notes: The table reports GMM parameter estimates and asymptotic standard errors (corrected for heteroskedasticity) in parentheses. The index A_t is constructed as the sum of the absolute standardized forecast errors associated with scheduled macro announcements on each day. We include U.S. announcements on: Non Farm Payroll, Durable Goods, Trade Balance, and Unemployment Claims, and German announcements on: the Current A/C, Employment, GDP, Import Prices, Industrial Production, M3, Manufacturing Orders, Manufacturing Output, Producer Prices, Retail Sales, the Trade Balance, Whole Sale Prices, and the Cost of Living Index. The standard errors reported for the variance ratios are computed from a Monte Carlo experiment with 5000 replications. $R_{\Delta p}^2(direct)$ and $R_{\Delta p}^2(indirect)$ are the fraction of the daily variance in prices attributed to news arrival via the direct and indirect channels.

Appendix Table 2
Variance Regressions (Daily Data)

Equation	Regressors				Diagnostics				
	Const.	A_t	A_t^{US}	A_t^G	R^2	SEE	Non-linear	Serial	Hetero
Price change	70.709 (12.271)	3.900 (0.793)			0.17	48.23	0.49	0.19 0.74	0.61 0.98
	69.404 (11.649)		6.033 (2.220)	3.367 (1.013)	0.18	48.27	0.57	0.21 0.81	0.64 0.99
Order flow	2.911 (0.732)	0.212 (0.050)			0.13	3.02	0.45	0.00 0.06	0.00 0.04
	2.881 (0.697)		0.258 (0.128)	0.201 (0.067)	0.13	3.03	0.81	0.00 0.05	0.00 0.04

Notes: The table reports OLS regression coefficients with standard errors in parenthesis. The dependent variable is the integrated variance of price changes (upper panel) and the integrated variance of order flow (lower panel). (In both cases the integrated variance is calculated using a 5-minute sampling frequency.) Regressors include; the number of US news arrivals A_t^{US} , German news arrivals, A_t^G , and total news arrivals A_t , all on day t . The Non-linear column presents the p-value of a chi-squared LM test for exclusion of the squared news term (or terms, in the case where country news arrivals are separated). The Serial column presents the p-value of a chi-squared LM test for first-order (top row) and fifth-order (bottom row) serial correlation in the residuals. The Hetero column presents the p-value of a chi-squared LM test for first-order (top row) and fifth-order (bottom row) ARCH in the residuals.

Appendix Table 3

Tests for State-Dependency (Intraday Data)

Non-linearity				
Equation	$\Theta(L)x_i n_i$	(p-value)	$\Theta(L)x_i A_i^*$	(p-value)
Δp^{ask}	113.277	(0.00)	13.044	(0.04)
Δp^{bid}	102.597	(0.00)	17.149	(0.01)
x	8.538	(0.20)	6.090	(0.41)

Heteroskedasticity				
	n_i	(p-value)	A_i	(p-value)
ε_t	797.285	(0.00)	4.552	(0.03)
ω_t^{ask}	5.544	(0.02)	1.266	(0.26)
ω_t^{bid}	3.207	(0.07)	0.489	(0.48)
v_t	6803.326	(0.00)	12.896	(0.00)

Notes: The upper panel reports Wald statistics and p-values for the null hypothesis of zero coefficients on the terms listed at the head of each column in models of the form:

$$z_i = \Theta(L)x_i + \Theta(L)x_i n_i + \Theta(L)x_i A_i^* + w_i .$$

All models were estimated by GMM allowing for heteroskedasticity and an MA(1) error structure in the case of $z_i = \Delta p_i^{ask}$ and $z_i = \Delta p_i^{bid}$. These estimates are then used to construct the Wald test. For the price change models, $\Theta(L) = \theta_1 L^{-4} + \theta_2 L^{-3} + \dots + \theta_6 L$ and for order flow, $\Theta(L) = \theta_1 L + \theta_2 L^2 + \dots + \theta_{10} L^{10}$. A_i^* and A_i are a dummy variables equal to one if there was a macro news arrival during the previous 15-mintues (i.e. 3 observation windows) or previous 5-mintues respectively. The lower panel reports the results of Glesjer (1969) tests for heteroskedasticity in the variance of each shock using the variables listed at the head of each column. The shock v_i is the innovation to the AR(10) model for order flow, while ε_i , ω_i^{ask} , and ω_i^{bid} are the shocks from

$$\Delta p_i^{ask} = \Theta(L)x_i + \varepsilon_i + \omega_i^{ask} - \omega_{i-1}^{ask}$$

$$\Delta p_i^{bid} = \Theta(L)x_i + \varepsilon_i + \omega_i^{bid} - \omega_{i-1}^{bid}$$

estimated jointly by GMM.

Appendix Table 4

Diagnostics for The Intraday Model

	<u>Statistic</u>	<u>P-value</u>					
Wald Test for $D(L, 0, 0) = D(L, \infty, 0)$	250.992	(<0.001)					
Wald Test for $D(L, 0, 1) = D(L, \infty, 1)$	16.507	(0.011)					
Wald Test for $D(L, 0, 1) = D(L, 0, 1)$	19.960	(0.003)					
Wald Test for $D(L, \infty, 1) = D(L, \infty, 0)$	16.034	(0.014)					
Test (i) for misspecification in $D(L, n_i, A_i^*)$	0.334	(0.999)					
Test (ii) for misspecification in $D(L, n_i, A_i^*)$	7.128	(0.309)					
Test (i) for misspecification in $\Sigma_\varepsilon(n_i, A_i)$	0.876	(0.349)					
Test (ii) for misspecification in $\Sigma_\varepsilon(n_i, A_i)$	0.218	(0.641)					
Test (i) for misspecification in $\Sigma_\omega(n_i, A_i)$	0.415	(0.520)					
Test (ii) for misspecification in $\Sigma_\omega(n_i, A_i)$	0.338	(0.561)					
Test (i) for misspecification in $\Sigma_v(n_i, A_i)$	2.522	(0.112)					
Test (ii) for misspecification in $\Sigma_v(n_i, A_i)$	0.028	(0.866)					
Squared Residual Autocorrelations (Std. Errs.)							
	Lag = 1	2	3	4	5	6	12
<u>Residual</u>							
$\varepsilon_i^2 / \Sigma_\varepsilon(n_i, A_i)$	0.090 (0.028)	0.038 (0.019)	0.047 (0.027)	0.161 (0.135)	0.028 (0.015)	0.058 (0.021)	0.019 (0.008)
$(\omega_i^{ask})^2 / \Sigma_\omega(n_i, A_i)$	0.068 (0.024)	0.015 (0.016)	0.015 (0.020)	-0.015 (0.011)	0.002 (0.010)	0.019 (0.012)	0.014 (0.011)
$(\omega_i^{bid})^2 / \Sigma_\omega(n_i, A_i)$	0.058 (0.024)	0.022 (0.010)	0.009 (0.012)	0.005 (0.012)	-0.003 (0.009)	0.010 (0.010)	0.023 (0.014)
$v_i^2 / \Sigma_v(n_i, A_i)$	0.068 (0.026)	0.032 (0.017)	0.039 (0.017)	0.033 (0.017)	0.000 (0.017)	0.022 (0.018)	-0.009 (0.007)
<p>Notes: The table reports diagnostic tests for the estimated intraday model shown in Table 3. The upper panel reports Wald tests for state-dependency in $D(L, n_i, A_i^*)$. The second panel reports LM-type tests for misspecification in the estimated $d_j(n_i, A_i^*)$ and $\Sigma_j(n_i, A_i)$ functions. Type (i) and (ii) tests consider misspecification with respect to trade intensity n_i and news arrival dummy A_i, respectively. (A_i^* and A_i are a dummy variables equal to one if there was a news arrival during the previous 15-mintues—i.e. 3 observation windows—or previous 5-mintues respectively.) The lower panel reports autocorrelations in the standardized squared shocks.</p>							

Appendix Table 5: Lower Bound on Percentage Price Variance From Order Flow Shocks

News Arrival					
	Horizon (minutes)				
	5	30	60	120	∞
<u>Trade intensity n</u>					
10	1.98 (1.20)	4.81 (3.00)	5.55 (3.58)	6.31 (4.33)	6.53 (4.61)
20	4.86 (2.04)	10.10 (4.17)	11.22 (4.76)	12.30 (5.46)	12.60 (5.70)
30	8.80 (3.23)	16.09 (5.51)	17.42 (6.05)	18.65 (6.63)	18.97 (6.81)
No News Arrival					
	Horizon (minutes)				
	5	30	60	120	∞
<u>Trade intensity n</u>					
10	0.84 (0.24)	3.01 (0.86)	3.85 (1.10)	4.92 (1.40)	5.28 (0.25)
20	1.75 (0.45)	4.55 (1.16)	5.27 (1.35)	6.02 (1.54)	6.23 (0.51)
30	2.69 (0.67)	5.75 (1.43)	6.39 (1.59)	6.98 (1.73)	7.15 (0.81)
P-values: No increase in variance ratio					
	Horizon (minutes)				
	5	30	60	120	∞
<u>Trade intensity n</u>					
10	0.04	0.08	0.10	0.12	0.00
20	0.00	0.01	0.01	0.01	0.00
30	0.00	0.00	0.00	0.00	0.00
Notes: Estimated from five-minute data (model in Table 3). The upper two panels report the estimated lower bound for the percentage of price variance over various horizons due to order flow shocks, given trade intensity n , with and without concurrent news arrival. This lower bound is estimated using only the price effects of interdealer order flow that is temporally prior to those price effects. (From Table 1, the median number of trades n per five-minute interval is 12 and the 75 th percentile is 21.) The standard errors reported in parenthesis in the upper panels, and p-values shown in the lower panel are calculated from 5000 Monte Carlo draws from the estimated asymptotic distribution of the GMM estimates.					