

How Is Macro News Transmitted to Exchange Rates?

March 21, 2006

Martin D. D. Evans¹
Georgetown University and NBER
Department of Economics
Washington DC 20057
Tel: (202) 338-2991
evansm1@georgetown.edu

Richard K. Lyons
U.C. Berkeley and NBER
Haas School of Business
Berkeley, CA 94720-1900
Tel: (510) 643-2027
lyons@haas.berkeley.edu

Abstract

Using data that avoid the stale quotes problem, we show that order flow contributes more to currency price changes following public news than at other times. This is inconsistent with news effects being common knowledge and therefore impounded in price directly. Roughly two-thirds of the total effect of macro news on the DM/\$ price is transmitted via order flow. With both the direct and indirect channels operating, macro news accounts for 36 percent of total price variance, an order of magnitude more than earlier estimates.

Keywords: Information, Microstructure, Announcements, Heterogeneity.

¹Respective 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, David Levine, Antonio Mello, Michael Melvin, Maureen O'Hara, Paul Pfleiderer, John Rogers, Andrew Rose, two anonymous referees, and seminar participants at the American Finance Association meeting, the NBER Microstructure meeting, U.C. Berkeley, Alberta, Michigan, London Business School, the New York Fed, Aix-Marseille, Virginia, Wharton, Insead, LSE, the Central Bank of Norway, and the Central Bank of Italy. Both authors thank the National Science Foundation for financial assistance, including funding for a web clearinghouse for micro-based FX research (see faculty.haas.berkeley.edu/lyons, georgetown.edu/faculty/evansm1).

1 Introduction

All textbook models of currency prices imply that public news determines prices directly: currency demand shifts are common knowledge and any related transactions play no role in causing the change. In microeconomic models of asset prices, transactions do affect prices causally (e.g., Glosten and Milgrom 1985, Kyle 1985). The causal role arises because transactions convey information that is not common knowledge. We test whether transactions transmit macroeconomic news to currency prices, and how this channel compares to the direct channel.

We examine the impact of macro news both intradaily and daily. We begin at the 5-minute frequency. Estimates of our intraday model using interdealer order flows show that while order flow contributes significantly to changing prices at all times, it contributes more to changing prices immediately after news arrival.² This is inconsistent with the textbook view that macro news effects are common knowledge and therefore impounded in prices without any order flow role. It suggests, instead, that macro news triggers trading that reveals dispersed information, which in turn affects currency prices.

Our daily analysis provides further evidence that trading on news reveals incremental information. The daily model distinguishes three sources of currency price variation. The first source mirrors traditional models – macro news that is impounded immediately and directly. The second source is the indirect effect of news on price via induced order flow. The third source is order flow that affects price but is unrelated to public news (possibly induced by banks’ changing risk tolerances, firms’ changing hedging demands, or individuals’ changing liquidity demands; see, e.g., Evans and Lyons 2002a). We find that all three sources of DM/\$ price variation are significant. Order flows vary considerably with macro news, with the result that roughly two-thirds of the effect of macro news on currency prices is transmitted via order flow, the remainder being the direct effect of news. This is consistent with the intraday finding that order flow is most important for determining currency prices during periods immediately following news arrival. With both the direct and indirect channels operating, we find that macro news accounts for 36 percent of total daily price variance. This is an order of magnitude more explanatory power than in previous studies (addressed below).

Though the literature on news and currency prices is long standing, it has not used quantities (order flow) to sort out the relationship. The literature has two branches: a first-moment branch that addresses the direction of price changes and a second-moment branch that addresses price volatility. A common finding of the first-moment branch is that directional price effects from scheduled macro announcements are difficult to detect at the daily frequency – they are swamped by other factors. Intraday event studies, such as Andersen et al. (2003), do find statistically significant effects, particularly for employment and money-supply

²Order flow is the cumulation over time of signed trades, where trades are signed according to whether the initiator is buying or selling (the marketmaker posting the quote is the non-initiating side). Order flow’s role in determining currency prices is documented by Payne (2003), Rime (2000), Evans and Lyons (2002a,b), and Evans (2002), among many others. Flows from individual end-user segments in currency markets are addressed in Lyons (2001), Froot and Ramadorai (2005), and Evans and Lyons (2005), among others. Order flow is similarly important for prices in bond markets, which share many informational and structural features with currency markets (see, e.g., Green 2004, Fleming 2003, and Brandt and Kavajecz 2004).

announcements.³ The second-moment branch focusing on news effects on volatility is partly a response to difficulty in finding news effects on return first moments.⁴ This work finds that announcements do indeed produce the largest price changes.

The two papers most closely related to our own are Green (2004) and Love and Payne (2004). Green focuses on the bond market and documents a significant increase in the informational role of trading following economic announcements. Information asymmetry increases following the release of public information in a way consistent with, for example, the skilled information processor models of Kim and Verrecchia (1994,1997); see also Kandel and Pearson (1995). The Love and Payne (2004) paper, concurrent with our own, uses order flow to sort out macro news effects on currency prices. They also find that a substantial share of public news is incorporated into prices via trading. Though there are many differences across the analyses, one key difference is exposure to the stale quotes problem: Love and Payne use data from limit-order trading, whereas we use data on bilateral direct transactions. If there is any staleness in existing limit orders when unambiguous news arrives, the first market orders that sweep out stale limit orders will be correlated with price change, but do not truly cause it (Carlson and Lo 2004 has some discussion of this causality concern). The data on bilateral direct transactions we use are free from this staleness.

Our paper departs from earlier work on news in two main ways: we exploit a much broader set of macro events and we employ a different identification strategy. The set of announcements used in past work is limited to scheduled announcements, which account for only about 10 percent of the news arrivals that appear on trading-desk news screens (Reuters Money Market Headline News). Our wider set gives us a more complete picture of the joint dynamics of currency (FX) prices and order flows, but requires that we account for the absence of complete data on ex-ante announcement expectations. For identification, like Rigobon and Sack (2004) we base our estimation strategy on state-dependent heteroskedasticity.⁵ Specifically, we identify the relative importance of direct and indirect news effects by allowing news to affect the variances of order flow and price differently. This approach does not require data on ex-ante expectations. Instead, it requires the weaker assumption that one can identify changes in the variance of macro information shocks. To ensure the robustness of our results, we model these variations in several different ways (in both the intraday and daily analysis).

Our finding that macro news accounts for more than one third of price variance helps to resolve a big puzzle in international finance – the news puzzle. The puzzle is that even the most comprehensive studies of news effects on currency prices account for less than 5 percent of total price variation. A good example

³See also, for example, Cornell (1982), Engel and Frankel (1984), Hakkio and Pearce (1985), Ito and Royley (1987), Hardouvelis (1988), Klein (1991), and Ederington and Lee (1995). For bond markets, see Fleming and Remolona (1997) and Balduzzi, Elton, and Green (2001).

⁴See, for example, Goodhart et al. (1993), DeGennaro and Shrieves (1997), Andersen and Bollerslev (1998), and Melvin and Yin (2000). For bond markets, see Fleming and Remolona (1999), Bollerslev, Cai, and Song (2000), and Huang, Cai, and Wang (2002).

⁵See the discussion in Rigobon and Sack (2004) 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.

at the daily frequency is Klein (1991). He regresses FX price changes on trade-balance news and finds that news explains about 40 percent of price changes on those days. This is an impressive finding. However, since trade balance news arrives monthly, roughly 95 percent of FX price variation is not included in the regression (20 of 21 trading days per month). Thus, an R^2 statistic of 0.4 implies that less than 3 percent of total price variation is accounted for. Andersen et al. (2003) also report impressive R^2 statistics within their event windows (in this case, intraday windows). But as they note (p. 50), summing the amount of time in all of their five-minute, post-event windows accounts for only 0.2 percent of their full sample period (e.g., roughly one five-minute interval per day). Under the conservative assumption that news arrival causes variance to increase by a factor of 10, their findings imply that news accounts for no more than 2 percent of the total price variation.⁶ We estimate the contribution of macro news to be an order of magnitude higher for two main reasons. First, our joint use of price and quantity data opens a new channel through which macro news can effect FX prices. Second, as noted, we consider the effects of a much broader set of macro news events. While scheduled announcements are certainly an important source of macro news, our estimates indicate that they do not encompass all the sources of macro news available to market participants.

The remainder of the paper is in four sections. Section 2 describes our data and presents some descriptive statistics. Section 3 presents the intraday analysis. Daily analysis is presented in Section 4. Section 5 concludes.

2 Data and Descriptive Statistics

Our order flow and price data are drawn from time-stamped, tick-by-tick transactions in the DM/\$ spot market over a four-month period, May 1 to August 31, 1996. The transactions are from the Reuters Dealing 2000-1 system that operates 24 hours a day, 7 days a week. Importantly, Dealing 2000-1 is a bilateral interdealer system on which a dealer requests a quote from another dealer, and when received, generally has only a few seconds to act before the quote is retracted. This type of data avoids the stale quote problem that can cloud inferences about causality when news arrives since, unlike limit orders, these quotes are always very short lived, and are generally not extended at moments of anticipated public news arrival. In 1996 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. Transactions between marketmakers accounted for about 75 percent of total trading in major spot markets at the time. This 75 percent breaks into two transaction types—direct (bilateral) and

⁶Security-return volatility is not constant over time (French and Roll 1986). Our daily-frequency example from Klein (1991) could include two adjustments in this respect: currency price 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. Andersen and Bollerslev (1998) report that Employment Report has the largest impact on the instantaneous variance, increasing it by a factor of 10. If all announcements had the same effect, and the within-event-window R^2 statistics were all one, news would still only account for 2 percent of the total exchange rate variation. In fact, the R^2 statistics in Andersen et al. (2003) are generally below 50 percent (Table 2), so the 2 percent figure is indeed an upper bound.

brokered (multilateral). Direct trading accounted for about 60 percent of trades between market-makers and brokered trading accounted for about 40 percent. (For more detail on this Reuters Dealing System see Lyons 2001 and Evans 2002; the latter includes details on data collection and statistical properties.) For every trade executed on D2000-1, our data set 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.

Our intraday analysis uses transaction prices, order flow and trade intensity measured over fixed intervals of five-minutes. We denote the last DM price for the purchase and sale of dollars in interval i as p_i^{ASK} and p_i^{BID} respectively. (With roughly 1 million transactions per day, the preceding transaction is only seconds before the end of each 5-minute interval during regular trading hours.) Interdealer order flow, x_i , is the difference during interval i between the number of trades initiated by dealers buying dollars and the number initiated by dealers selling dollars.⁷ Similarly, we measure trade intensity, n_i , by the unsigned number of interdealer transactions per minute during interval i . Daily versions of our series are denoted with subscript t for the daily price p_t , we use the last DM price for the purchase of dollars before 5 pm BST (British Summer Time).⁸ Daily order flow, x_t , is the same as five-minute order flow x_i save that it spans the time difference between 5 pm on days $t - 1$ and t . Trading intensity on day t , n_t , is defined as the number of transactions over the same daily interval.

The primary source of our news data is the Reuters Money Market Headline News screen (archived by Olsen Associates). These screens are standard equipment on FX trading desks and are used for high frequency monitoring by non-marketmaker participants as well. Reuters collects news reports from approximately 150 bureaus around the world. Each report must be approved by an economics editor at Reuters before it appears as a news item on the Headline screens. The presence of this editorial process means that all the news items in our data set were viewed as containing news-worthy economic information. At the same time, competition between Reuters, Bloomberg and Dow Jones insures that editorial decisions minimize publication delay. We impose a further layer of editorial screening by excluding from our data set news items of the following four types: (i) reports of upcoming known holidays, (ii) reports that a scheduled data release would take place (e.g., “Monthly employment report due out tomorrow”), (iii) duplicate reports (the same news is repeated with a slight change in wording), and (iv) reports referring to the DM/\$ price or market. The four filters exclude less than 10 percent of news arrivals. The first three filters are intended to distill information that is truly incremental.

A number of other factors give us confidence that our analysis is not significantly exposed to feedback from the DM/\$ market to macro news flow. The potential here is that increased volatility in the DM/\$ price creates incentives for reporters to initiate news items to explain it, which are then posted to the

⁷In direct trading between marketmakers, order sizes are standardized, so variation in size is much smaller than variation in the size of individual trades between marketmakers and their end-user customers. 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, which we use below to interpret the estimated coefficients.

⁸Using prices from buyer-initiated transactions eliminates return reversals from prices bouncing randomly from bid to ask.

Headline screen. Our fourth filter helps to protect against this form of endogeneity insofar as the news item makes reference to the DM/\$ market. The well-defined editorial process described also helps protect against spurious news creation. Perhaps most important, the Headline screen is used by traders in many markets (money markets, bond markets, currency markets, and others), so the audience is much wider than just the DM/\$ market. We find the hypothesis of feedback to news flow patently strained when it comes to our analysis at the five-minute frequency: the intermittency of arrivals shown in Table 1 makes this kind of ultra-high-frequency feedback hard to imagine.

Table 1: Macro News Sample

Date	Time	News
5/1/1996	13:05:22	MARCH U.S. LEADING INDICATORS SHOW ECONOMY EASING
5/1/1996	14:00:50	U.S. MARCH CONSTRUCTION SPENDING ROSE 3.1 PCT
5/1/1996	14:10:14	MARCH U.S. CONSTRUCTION SPENDING REBOUNDS STRONGLY
5/2/1996	6:05:18	GERMAN MARCH IMPORT PRICES CLIMB 0.3 PCT M/M
5/2/1996	8:33:10	BUNDESBANK DOES NOT PLAN NEWS CONFERENCE TODAY
5/2/1996	9:48:20	GERMAN CALL MONEY FALLS BACK TO 3.30/40 PCT
5/2/1996	10:50:08	BUNDESBANK LEAVES INTEREST RATES UNCHANGED
5/2/1996	10:51:56	GERMAN MARCH INDUSTRIAL OUTPUT DATA DUE 1130 GMT
5/2/1996	12:30:40	U.S. JOBLESS CLAIMS FELL IN LATEST WEEK
5/2/1996	12:31:00	U.S. Q1 1996 REAL GDP ROSE 2.8 PCT
5/2/1996	14:00:30	U.S. MARCH FACTORY ORDERS ROSE 1.5 PCT
5/2/1996	15:01:54	U.S. Q1 GDP SURGE SEEN JUST A BLIP IN MODEST TREND
5/2/1996	15:10:32	GERMAN EMPLOYER TO OPPOSE ANTI-WAGE DUMPING LAW
5/3/1996	9:56:32	GERMAN CALL MONEY EASES SLIGHTLY AHEAD OF WEEKEND
5/3/1996	12:30:38	U.S. APRIL NON-FARM PAYROLLS ROSE 2,000
5/3/1996	12:31:00	U.S. MARCH PERSONAL INCOME ROSE 0.5 PCT
5/3/1996	12:38:44	U.S. JOBLESS RATE LOWER BUT LABOR MARKET LOOKS WEAK
5/3/1996	13:42:16	MARCH US INCOME DATA SHOW MODEST GROWTH, INFLATION
5/3/1996	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). Time is measured as British Summer Time (BST).

For further perspective on the news data, Table 1 lists the filtered news items over the first three days of the sample. These items clearly extend well beyond scheduled releases of macroeconomic data.⁹ If we were to focus only on the small subset of data releases that are scheduled, as we do when checking the robustness of our results, there would be many trading days in our sample without a single news event. The use of our broader set shows that in fact these were days on which macro news did arrive. Table 1

⁹A wider set of news types is also considered by Eddelbuttel and McCurdy (1998), though their focus is 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.

illustrates another feature of our data: although we filter out duplicate news items, we retain items that *interpret* previous information. For example, we retain the third item listed in Table 1 because it interprets, rather than simply restates, the information on construction spending contained in the second item. Does such an interpretation represent news? Clearly the 3.1% increase in construction spending could have been interpreted by *some* as representing a strong rebound, but it seems far-fetched to assume that *everyone* subscribing to Reuters held this view and recognized the unanimity of opinion. When there is anything short of a unanimous interpretation of a data release, a subsequent news item providing interpretation will contain new information to at least some agents. Our prior is that data releases rarely (if ever) meet this unanimity requirement and so we retain the interpretive items in our data set.¹⁰

Importantly, the estimation strategy we adopt in both our intraday and daily analysis does not require that every news item is equally important. As we detail below, all we require is that the news data can be used to identify variations in the flow of macro news hitting the FX market. For this purpose we construct several different measures of macro news flow: one based on the arrival rates of just US news items, one based on German items, and one based on the arrival of both US and German items. We also construct measures from the small subset of releases that are scheduled. For this subset we can combine the Reuters data with survey data on ex-ante expectations (provided by Money Market Service), allowing us to construct estimates of macro news flow from estimates of unexpected components.¹¹ We use these different measures of macro news flow to check the robustness of our estimation results. In particular, since the arrival of scheduled news is by definition immune to possible feedback from FX price volatility to the arrival of unscheduled news, comparing results using all news versus scheduled news provides a test for the presence of feedback.

Table 2 presents descriptive statistics for the variables used in intraday and daily analyses. The upper rows of panel A report sample statistics for the daily change in FX prices, Δp_t , and the level of interdealer order flow x_t . There is no detectable serial correlation in either variable at the daily frequency: The estimated first order autocorrelation in the Δp_t and x_t series are 0.015 and -0.035, and are both statistically insignificant. The remaining rows in panel A report statistics on four of our measures of macro news flow. A_t^{US} and A_t^{GM} respectively denote the number of US and German news items appearing on the Reuters Headline screen between 5:01 pm BST on day $t - 1$ and 5 pm BST on day t . A_t is the daily arrival rate of all news computed as $A_t^{\text{US}} + A_t^{\text{GM}}$. A_t^{S} denotes the arrival rate for the subset of scheduled news, defined as the number of scheduled releases between 5:01 pm BST on day $t - 1$ and 5 pm BST on day t concerning the following variables for the US: Non Farm Payroll, Durable Goods, Trade Balance, and Unemployment Claims, and for Germany: Current Account, Employment, GDP, Import Prices, Industrial Production, M3, Manufacturing Output and Orders, Retail Sales, Trade Balance, and the Cost of Living. As the table shows, the median arrival rate

¹⁰Note, also, that we retained the eighth item listed in Figure 1 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. There was no record of the news item containing the German Industrial Production data in the archive created from the news screens by Olsen Associates – we do not know why.

¹¹Event studies in FX, such as Andersen et al. (2003), regress returns (price changes) on innovations for individual data types (e.g. Unemployment Claims) computed from the MMS data. We cannot study the impact of scheduled news at this disaggregate level because our four-month sample contains too few scheduled releases for each data type.

Table 2: Sample Statistics

	Min	5%	25%	50%	75%	95%	Max	Std.	Skew.	Kurt.
A: Daily Data										
Δp_t	-20.7	-11.9	-3.8	0.3	3.4	6.9	12.4	5.9	-0.8	3.8
x_t	-449.0	-308.0	-61.0	8.0	91.0	186.0	339.0	136.4	-0.6	4.5
A_t^{US}	0	0	1	2	5	7	9	2.2	0.7	2.6
A_t^{GM}	0	2	6	8	12	18	22	5.1	0.5	3.0
A_t^{ALL}	0	2	9	11	15	21	27	6.0	0.2	2.6
A_t^{S}	0	0	0	1	2	4	7	1.5	1.6	5.4
B: Intraday Data										
Δp_i	-7.9	-1.4	-0.3	0.0	0.3	1.3	5.0	0.8	-0.2	7.4
x_i	-72.0	-9.0	-2.0	0.0	3.0	9.0	69.0	5.6	0.1	12.6
n_i	2.0	2.0	6.0	12.0	21.0	44.0	212.0	15.5	3.4	23.2
Autocorrelations										
Lag =	1	2	3	4	5	6	12	18	24	
Δp_i	-0.31	-0.00	-0.00	-0.00	-0.00	0.00	0.00	0.02	0.00	
	(<.01)	(0.35)	(0.76)	(0.79)	(0.68)	(0.23)	(0.69)	(0.06)	(0.64)	
x_i	0.23	0.10	0.09	0.08	0.06	0.06	0.03	0.02	0.00	
	(<.01)	(<.01)	(<.01)	(<.01)	(<.01)	(<.01)	(<.01)	(0.01)	(0.65)	

Notes: The sample is May 1 to August 31, 1996. Δp_t is 1000 times the change in the last DM purchase price for dollars between 5:00 pm on day t and day $t-1$. x_t is the total interdealer order flow over the same time interval. A_t^{US} and A_t^{GM} are respectively the number of macro news arrivals observed on the Reuters Money Market Headline News screen relating to the US and Germany between 5:00 pm on day t and day $t-1$. A_t^{ALL} denotes the total number of news items, $A_t^{\text{US}} + A_t^{\text{GM}}$, and A_t^{S} is the total number of scheduled news items arriving over the same time interval. Scheduled announcements include US Non Farm Payroll, Durable Goods, Trade Balance, and Unemployment Claims and the German Current Account, Employment, GDP, Import Prices, Industrial Production, M3, Manufacturing Output and Orders, Retail Sales, Trade Balance, and the Cost of Living. In panel B, Δp_i and x_i are the price changes (DM purchase price for dollars) and order flows during 5-minute interval i , and n_i is the total number of trades. Autocorrelations are computed by GMM as in Evans (2002) and the p-values reported in parenthesis are calculated from Wald tests of the null hypothesis of a zero correlation (allowing for conditional heteroskedasticity).

for German news is four times the rate for US news. Notice also that the arrival rates for all news are considerably higher than the rate for scheduled news.

Panel B of Table 2 presents descriptive statistics for prices, order flow, and trade intensity measured at the 5-minute frequency. The sample statistics for Δp_i^{ASK} and Δp_i^{BID} are almost identical, so we only report those for Δp_i^{ASK} . One noteworthy feature of these statistics concerns the distribution of trade intensity, n_i . While the median trade intensity in our sample is 12 trades per minute, the distribution for n_i indicates that the pace of trading is occasionally much higher. Evans (2002) shows that some of the variations

in trade intensity can be related to the shift from predominantly Asian-based to US-based dealers as the trading day progresses. However, on a particular day, variations in trade intensity can differ significantly from this “seasonal” pattern. From the lower portion of panel B, we see a sharp difference from the daily frequency statistics: both price changes and order flows are serially correlated at high (intraday) frequencies. Transaction price changes (all are ask prices) display significant negative autocorrelation, but only at lag one, while order flow appears serially correlated at up to 18 lags.

We track the arrival of news at the 5-minute frequency with dummy variables. The dummy variable A_i takes the value of one if either a US or German news item appears on the Reuters screen during interval i . At least one news arrival occurs in 515 out of the 11,473 five-minute windows in our four-month sample. We use this dummy-variable approach in the five-minute data because there are few instances of more than one news arrival during a single five-minute observation window (in 29/515 there were two arrivals and in 4/515 there were three, numbers that proved insufficient to get mileage from a multi-valued dummy). We also make use of analogous dummies for all German news, all US news, and all scheduled US news; denoted respectively by A_i^{GM} , A_i^{US} and A_i^{s} .

3 Intraday Analysis

Our intraday analysis is based on a model for the joint dynamics of FX prices and order flows estimated at the 5-minute frequency. Information is impounded into FX prices via two channels. The first is the direct channel through which the arrival of new common-knowledge information leads marketmakers to change the FX prices they quote. The transmission of information into FX prices via this channel is direct and instantaneous. The second channel, the indirect channel, operates via order flow. In this case the arrival of information is first manifest in the trading decisions of individuals because the information is dispersed. Once marketmakers observe the ensuing order flow, they adjust their FX quotes to reflect the new information embedded in the pattern of trading. Thus, order flow is the medium by which dispersed information becomes embedded into FX prices.

Our intraday analysis will focus on the relative importance of the direct and indirect information channels in the period immediately following the arrival of news. The motivation for this focus is straightforward: If macro news primarily comprises new common-knowledge information, as is traditionally assumed, we should find evidence that the direct channel accounts for most of the FX price variation over intervals that include the news arrival. Conversely, if the arrival of macro news triggers revelation of dispersed information, possibly reflecting diverse views about price implications, we should find that the indirect channel dominates. We will quantify the relative importance of the direct and indirect channels from a decomposition of the variance in FX price changes.

3.1 The Model

Our intraday model extends the empirical model in Evans (2002) to account for the effects of news arrivals. At the heart of the model are the following equations:

$$\Delta p_i = B(L)\xi_i + \varepsilon_i, \tag{1}$$

$$y_i = C_y(L)\xi_i, \tag{2}$$

where Δp_i is the change in the spot price of FX between the end of periods $i-1$ and i , and y_i is the order flow initiated by end-users during period i . (The relationship between this end-user flow y_i and inter-dealer flow x_i is addressed below.) Equation (1) shows how prices respond to two types of news: common knowledge news shocks ε_i , and dispersed information shocks, ξ_i . We assume that these shocks are mutually orthogonal and serially uncorrelated (addressed below). The ε_i shocks represent unambiguous price-relevant news that is simultaneously observed by everyone and so are impounded fully and instantaneously into the price of FX. Dispersed information shocks represent, in aggregate, the bits of information contained in the trades of individual agents. This information is first manifested in the order flow, y_i , and then subsequently impounded in price. End-user order flow is the difference between the purchase and sales of dollars initiated by end-users at dealer FX quotes. The dynamic responses of prices and order flow to these dispersed information shocks are determined by the lag polynomials $B(L)$ and $C_y(L)$.

Three features of our specification deserve note. First, equation (1) describes the dynamics of transactions prices, p_i , defined as the market-wide average price at which actual transactions take place at time i . We will describe the link between this p_i and actual transactions below. Second, the assumed orthogonality between the common-knowledge and dispersed information news shock implies that common-knowledge news has no effect on order flow. This assumption has 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). Intuitively, any revision in price due to common-knowledge news should establish a new market-clearing price that does not systematically favor subsequent imbalances of sell orders over buy orders, or vice versa. For example, 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 unbiased. (Notice that this has nothing to do with the behavior of unsigned trading volume; our model does not restrict how common-knowledge news affects volume through, say, portfolio rebalancing.) The third feature concerns the dynamics of end-user order flow y_i : We assume that end-users' demand for foreign currency is imperfectly elastic, so any imbalance in order flow (i.e., $y_i \neq 0$) requires price adjustment to achieve market clearing. Consequently, all order flow is, at least temporarily, price relevant.¹² Under rational expectations, this information is summarized in current and past dispersed

¹²Our elasticity assumption does not imply that shocks to order flow necessarily have permanent price effects. It is possible that some shocks to order flow only affect prices while the associated inventory imbalance is being spread among dealers (see

information shocks, but remains unrelated to common-knowledge news shocks, as shown in equation (2).

Equations (1) and (2) allow us to identify three channels through which the arrival of macro news may affect the dynamics of price and order flows. First, when the macro announcement contains a common-knowledge component, it will affect prices instantaneously via the ε_i shock. This direct channel will be operable when everyone agrees on the price-implications of the announcement. Second, when a macro announcement is viewed by different agents as having different price implications, its effects on prices and order flow will manifest via the ξ_i shocks: Although everyone observes the same announcement, different views about the mapping from macro data to FX prices represent dispersed information that is relevant for equilibrium prices. Third, the arrival of a macro announcement can affect the *process* through which dispersed news is impounded into prices, by which we mean the lag polynomials. We allow for this by allowing $B(L)$ and $C_y(L)$ to vary with the arrival of new announcements.

3.1.1 Empirical Specification

Estimation of our intraday model is complicated by two factors: First, our data are on market-wide order flow between dealers, x_i , rather than the end-user order flows y_i . We must be careful to distinguish these different order flows if we are to account for the temporal impact of dispersed information. Second, our model needs to accommodate forms of state-dependency beyond the arrival of macro news. We shall deal with these complications in turn.

Prices in the data set come in two forms. If a dealer initiating a transaction buys dollars, the transaction price equals the ask quote in DMs per dollar offered by the other marketmaker. We refer to this as the DM purchase price for dollars, p^{ASK} . If the dealer initiating a transaction sells dollars, the transaction price will equal the bid quote given by the other dealer. We refer to this as the DM sale price for dollars, p^{BID} . Evans (2002) finds evidence that lack of transparency in direct dealer trading allows for an equilibrium price distribution, as opposed to a strict law of one price. To formalize this idea, our intraday model assumes that equilibrium in the market at a point in time is described by a distribution of purchase prices and a distribution of sales prices.

Let p_i^{ASK} and p_i^{BID} denote observed prices drawn randomly from the respective distributions of purchase and sales prices at time i . These observed prices are related to the average transaction price, p_i , defined in (1), by:

$$p_i^o = p_i + \eta_i^o, \tag{3}$$

for $o = \{\text{ASK}, \text{BID}\}$. η_i^{ASK} and η_i^{BID} are idiosyncratic shocks that identify the degree to which observed prices differ from the market-wide average. Their size depends on the identity of the dealers whose prices we observe. We assume that observed prices are drawn randomly and independently from the cross-sectional

Cao, Evans and Lyons 2006). In this special case, some of the individual coefficients in $B(L)$ will differ from zero, but their sum will equal zero.

distributions of purchase and sale prices every period, so that η_i^{ASK} and η_i^{BID} are serially uncorrelated and independently distributed.

The second complication arises from the distinction between the interdealer and end-user order flows. The order flow measure in our data set is derived from trades initiated between dealers. These trades are temporally downstream from the trades initiated by end-users against dealer quotes. As a result, it is possible for a dispersed information shock ξ_i to affect prices and end-user order flows before it shows up in interdealer order flow: Dealers may adjust their price in the face of an end-user order induced by ξ_i *before* initiating trades in the interdealer market for risk sharing or speculative motives. Thus, price changes may appear temporarily prior to changes in interdealer order flow even though they represent a response to earlier end-user order flow. We allow for this possibility by assuming that the interdealer order flow we measure is a distributed lag of end-user order flow:

$$x_i = C_x(L)y_{i-m}, \quad (4)$$

where, again, $C_x(L)$ is a polynomial in the lag operator. In this specification, it takes at least m periods before imbalances in end-user orders for FX show up in interdealer order flow (where m may be zero).

Combining (4) with (1) and (2), we can now represent the dynamics of prices and interdealer order flow by:

$$\Delta p_i = D(L)x_i + \varepsilon_i, \quad (5)$$

$$x_i = C(L)\xi_{i-m} \quad (6)$$

where $D(L) = B(L)L^{-m}C(L)^{-1}$ and $C(L) = C_x(L)C_y(L)$. Although the polynomial $D(L)$ may take many forms depending on the dynamic responses of price and interdealer order flow to dispersed information shocks, in general it will include both negative and positive powers of L (corresponding to leads and lags of x_i) when $m > 0$. Our model estimates are based on a sixth-order specification for $D(L)$ (shown below) that links Δp_i to interdealer order flows from x_{i+4} to x_{i-1} . This specification is supported by a series of diagnostic tests reported in Evans (2002). It implies that a dispersed information shock may impact end-user orders and prices up to 20 minutes before it affects interdealer order flow (i.e., $m = 4$). Similarly, we specify the form of $C(L)$ so that the time series properties implied by (6) match those in the data. As in Evans (2002), we find that interdealer order flow is well characterized by an AR(10) process, so we specify $C(L)$ as $(1 - \sum_{j=1}^{10} c_j L^j)^{-1}$.

Finally, we incorporate the effects of macro news. We treat the arrival of news as changing the state of the market. Following Evans (2002), we also allow the dynamics of prices and order flow to vary with trading intensity. Including trading intensity as a state variable is important for accommodating the pronounced time dependence in volatility documented by Andersen and Bollerslev (1998). Let S_i denote the state of the market in period i . We assume that S_i depends on trading intensity in period i , n_i , and the arrival of news during the past three periods, A_i , A_{i-1} and A_{i-2} . (Recall that the dummy variable A_i equals one if a macro

news arrives during period i .) We incorporate state-dependency into the price and order flow dynamics via the polynomial $D(L)$, and the error variances. Specifically, $D(L)$ is replaced by $D(L, S)$, a state-dependent sixth order polynomial:

$$D(L, S) = d_1(n, \bar{A})L^{-4} + d_2(n, \bar{A})L^{-3} + \dots + d_5(n, \bar{A}) + d_6(n, \bar{A})L. \quad (7)$$

where $\bar{A}_i \equiv \max\{A_i, A_{i-1}, A_{i-2}\}$ with state-dependent coefficients $d_j(., .)$. Thus, $d_6(n, 1)$ is the coefficient on lagged order flow x_i when trade intensity equals n and news arrived in the past 15 minutes. We also allow for state-dependence in the error variances, $Var(\varepsilon_i|S_i) = \Omega_\varepsilon(n_i, A_i)$, $Var(\xi_i|S_i) = \Omega_\xi(n_i, A_i)$, and $Var(\eta_i^{\text{ASK}}|S_i) = Var(\eta_i^{\text{BID}}|S_i) = \Omega_\eta(n_i, A_i)$. State-dependence in the coefficients and variances is modeled as:

$$d_j(n, \bar{A}) = \underline{d}_j(\bar{A})e^{(-n/100)} + \bar{d}_j(\bar{A})[1 - e^{(-n/100)}], \quad (8)$$

$$\Omega_j(n, A) = \underline{\omega}_j(A)e^{(-n/100)} + \bar{\omega}_j(A)[1 - e^{(-n/100)}], \quad (9)$$

where $\underline{d}_j(0)$, $\bar{d}_j(0)$, $\underline{\omega}_j(0)$, and $\bar{\omega}_j(0)$ are the parameters to be estimated for observations without a news arrival, and $\underline{d}_j(1)$, $\bar{d}_j(1)$, $\underline{\omega}_j(1)$, and $\bar{\omega}_j(1)$ when there is a news arrival. These functional forms make $d_j(.)$ and $\Omega_j(.)$ smooth monotonic functions of the transaction rate and are similar to the transition functions used in nonlinear time series models (Potter 1999). They bound the coefficients between $\underline{d}_j(\bar{A})$ and $\bar{d}_j(\bar{A})$, and the variances between $\underline{\omega}_j(A)$ and $\bar{\omega}_j(A)$ as the transaction rate varies between 0 and ∞ .

Several aspects of our specification for state-dependency deserve comment. First, while specialized with respect to variations in trading intensity, the functional forms in (7) - (9) do not appear unduly restrictive when we subject our model to specification tests below. Second, there is no evidence that variations in trading intensity or the arrival of news affect the dynamics of order flow via $C(L)$. Thus, we do not incorporate state-dependency in this polynomial to avoid an unnecessary proliferation in parameters. Third, our specification places minimal restrictions on how the arrival of news affects the error variances and the link between order flow and price dynamics. Importantly, we do not restrict how the coefficients in $D(L, S)$ or the error variances change following the arrival of news. Consequently, our specification does not impose a prior about how the arrival of macro news affects the relative importance of the direct and indirect information transmission channels. Finally, our specification makes no distinction between the arrival of US news, German news, scheduled news or unscheduled news; A_i equals one when any news arrives during period i . We recognize that this assumption may be too restrictive. For example, it is possible that the information transmission process following the arrival of scheduled US news differs from that following other news items. Below we investigate the adequacy of this assumption with a series of specification tests.

3.1.2 Estimation

The model is estimated using the Generalized Method of Moments technique developed in Evans (2002). The moment conditions used to estimate the parameters of the order flow process are

$$0 = E[\xi_i \otimes z_i^x], \quad (10a)$$

$$0 = E[\{\xi_i^2 - \Omega_\xi(S_i)\} \otimes z_i^x], \quad (10b)$$

where $\xi_i = x_{i+4} - \sum_{j=1}^{10} c_j x_{i+4-j}$ and $\Omega_\xi(S_i)$ is the conditional variance of ξ_i specified in (9). (Hereafter, we use S_i rather than n_i and A_i as the argument of the error variances, $\Omega(\cdot)$.) If the order flow process is correctly specified, a dispersed information shock ξ in period i should be uncorrelated with interdealer order flow x in periods $i+3$ and earlier. Similarly, the difference between ξ_i^2 and the conditional variance should be uncorrelated with current or past trade intensity and order flows. We employ $\{x_{i+3}, \dots, x_{i-6}\}$ and four lagged values of ξ_i as elements of the instrument vector z_i^x in (10a). In (10b) the instrument vector contains a constant, $e^{(-n_i/100)}$ and A_i . With this choice of instruments, equations (10a) and (10b) represent 17 moment restrictions on 14 parameters ($\{c_j\}_{j=1}^{10}, \underline{\omega}_\xi(0), \underline{\omega}_\xi(1), \bar{\omega}_\xi(0)$ and $\bar{\omega}_\xi(1)$).

Parameters of the price process are computed from moments using the bivariate process for purchase and sales prices, Δp_i^{ask} and Δp_i^{bid} . Combining (3) with (5) and our specification for $D(L, S)$ gives:

$$\Delta p_i^o = \sum_{j=1}^6 \underline{d}_j(\bar{A}_i) x_{i+5-j} + \sum_{j=1}^6 [\underline{d}_j(\bar{A}_i) - \bar{d}_j(\bar{A}_i)] e^{-n_i/100} x_{i+5-j} + u_i^o$$

where $u_i^o \equiv \varepsilon_i + \eta_i^o - \eta_{i-1}^o$ for $o = \{\text{ASK, BID}\}$. This equation describes the state-dependent relation between actual transactions prices and interdealer order flow implied by our model. Notice that the composite error term u_i^o follows an MA(1) process and that $Cov(u_i^{\text{ASK}}, u_i^{\text{BID}}) = \Omega_\varepsilon(n_i, A_i)$. We account for this error structure in the moment conditions used to estimate the parameters of the price process:

$$0 = E[u_i^o \otimes z_i^p] \quad (11a)$$

$$0 = E[\{(u_i^o)^2 - \Omega_\varepsilon(S_i) - \Omega_\eta(S_i) - \Omega_\eta(S_{i-1})\} \otimes z_i^p] \quad (11b)$$

$$0 = E[\{u_i^o u_i^\emptyset - \Omega_\varepsilon(S_i)\} \otimes z_i^p] \quad (11c)$$

$$0 = E[\{u_i^o u_{i-1}^\emptyset + \Omega_\eta(S_{i-1})\} \otimes z_i^p] \quad (11d)$$

$$0 = E[u_i^o u_{i-1}^\emptyset \otimes z_i^p] \quad (11e)$$

$$0 = E[u_i^o u_{i-2}^\emptyset \otimes z_i^p] \quad (11f)$$

$$0 = E[u_i^o u_{i-2}^\emptyset \otimes z_i^p] \quad (11g)$$

for $o, \emptyset = \{\text{ASK, BID}\}$ and $\emptyset \neq o$. The moment restriction in (11a) exploits the assumed orthogonality between the instruments, z_i^p , and both the common knowledge news and idiosyncratic shocks. The other restrictions in (11) are derived from the moving average structure of the composite error. In particular, (11b) and (11c) focus

on the variance of $\{u_i^{\text{ASK}}, u_i^{\text{BID}}\}$, while (11d) - (11g) focus on the autocovariance. For example, in (11f) and (11g) we exploit the fact that under an MA(1) process, all the autocorrelations in the composite errors at lag 2 are zero. We use $\{x_{i+j}, e^{(-n_i/100)}x_{i+j}, \bar{A}_i e^{(-n_i/100)}x_{i+j}\}_{j=-1}^4$ as instruments in (11a), $\{1, e^{(-n_i/100)}, A_i\}$ in (11b) - (11d), and a constant in (11e) - (11f). This instrument choice gives us 66 moment restrictions on the 32 parameters of the prices process ($\{\underline{d}_j(0), \underline{d}_j(1), \bar{d}_j(0), \bar{d}_j(1)\}_{j=1}^6$ and $\{\underline{\omega}_j(0), \underline{\omega}_j(1), \bar{\omega}_j(0), \bar{\omega}_j(1)\}_{j=\varepsilon, \eta}$). Thus, the complete model contains 83 moment restrictions on 46 parameters.

In standard time series applications, GMM estimates of the parameter vector θ are found by minimizing a quadratic form constructed from the sample analogues of the moment conditions implied by the model. In this application, estimation is complicated by the fact that the gap between successive purchases and/or sales occasionally span many minutes. In these cases there is no record of an FX purchase and/or sale in the observation interval. For the purpose of computing our estimates, we designate the price, and order flow observations from these periods as “missing” and construct sample moments without these observations. Specifically, let $E[m_{i,j}(\theta)] = 0$ denote condition j among the moment conditions shown in (10) and (11) and let $\Lambda = \{i_1, \dots, i_T\}$ be the set of observations for which none of the elements in $m_{i,j}(\cdot)$ for all j is “missing”. We compute the sample analogue to condition j as $\bar{m}_j(\theta) = T^{-1} \sum_{\Lambda} m_{i,j}(\cdot)$. The GMM estimates of θ are then found by minimizing:

$$Q(\theta) = \bar{m}(\theta)' W^{-1} \bar{m}(\theta), \quad (12)$$

where $\bar{m}(\theta) = [\bar{m}_1(\theta), \bar{m}_2(\theta), \dots]'$. We follow the standard practice of first setting the weighting matrix W equal to the identity to obtain consistent estimates of θ . These estimates, $\tilde{\theta}$, then are used to compute a consistent estimate of the optimal weighting matrix, \tilde{W} . We construct \tilde{W} using the Newey and West (1987) estimator for the covariance of $m_{i,j}(\theta)$ incorporating a correction for MA(1) serial correlation. This estimate of the covariance matrix allows for the fact, documented below, that the model fails to completely account for the heteroskedasticity in prices and order flow. The GMM estimates, $\hat{\theta}$, are found by minimizing (12) with $W = \tilde{W}$. The asymptotic covariance matrix of the resulting estimates is $\hat{V} = [\hat{G} \tilde{W}^{-1} \hat{G}']^{-1}$ where $\hat{G} = \partial \bar{m}(\hat{\theta}) / \partial \theta'$.

We examine the performance of our estimated model with a series of diagnostic tests. In particular, we use a chi-squared test to examine the validity of an auxiliary set of moment conditions implied by our model but not used in estimation. Let $\bar{m}_{\text{II}}(\theta)$ denote a vector of K_{II} sample moments, comprising the K_{I} moments used to find the GMM estimates, and $K_{\text{II}} - K_{\text{I}}$ auxiliary moment conditions implied by the model. Following Hayashi (2000), we construct the test statistic by first finding the GMM estimates of θ , denoted $\hat{\theta}_{\text{II}}$, from the set of K_{II} moments. These estimates are found with the two-step procedure described above using the Newey and West estimator from the first step to construct the weighting matrix, \tilde{W}_{II} . Next, we construct the submatrix of \tilde{W}_{II} corresponding to the original K_{I} moments, \tilde{W}_{I} . We then find an alternative set of GMM estimates, $\hat{\theta}_{\text{I}}$, by minimizing (12) with $W = \tilde{W}_{\text{I}}$. Finally, we form the test statistic

$$C \equiv T \bar{m}_{\text{II}}(\hat{\theta}_{\text{II}})' \tilde{W}_{\text{II}}^{-1} \bar{m}_{\text{II}}(\hat{\theta}_{\text{II}}) - T \bar{m}(\hat{\theta}_{\text{I}})' \tilde{W}_{\text{I}}^{-1} \bar{m}(\hat{\theta}_{\text{I}}). \quad (13)$$

where T denotes the number of “non-missing” elements used to construct $\bar{m}_{11}(\theta)$. Under the null hypothesis that the auxiliary moment conditions are satisfied, the C statistic has an asymptotic chi-squared distribution with $K_{II} - K_I$ degrees of freedom. We use this test below to examine the adequacy of our specification for the state-dependent coefficients and error variances.

3.1.3 Model Estimates

Table 3 presents GMM estimates of the intraday model. In specifications where all the variance parameters were left unrestricted, the estimates of $\underline{\omega}_\varepsilon(A)$, $\underline{\omega}_\xi(A)$, and $\bar{\omega}_\eta(A)$ were very close to zero (i.e. < 0.0001), so the table reports estimates where these parameters are restricted to zero. With these restrictions imposed, there are 40 parameters to be estimated from a total of 64 moment restrictions, so our estimates are derived from a model with 24 over-identifying restrictions. The Hansen (1982) J -statistic computed from our GMM estimates is 55.722 which implies a p-value of 0.092 for the null of a correctly specified model.

Panel A of Table 3 reports the parameters for the state-dependent order flow polynomial, $D(L, S)$. A comparison of the estimates in rows (i) and (ii) and rows (iii) and (iv) shows that trade intensity has differing effects on the price-impact of order flow depending on the arrival of news. This is most easily seen in the right hand column where we report the sum of the coefficients in different market states. These estimates have two noteworthy features. First, the long run impact of order flow on prices is much larger when trading intensity is high ($\sum_j \underline{d}_j(\cdot) < \sum_j \bar{d}_j(\cdot)$). Second, controlling for trading intensity, the arrival of news slightly reduces the long-run impact of order flow ($\sum_j \underline{d}_j(\bar{A} = 1) < \sum_j \underline{d}_j(\bar{A} = 0)$, except at the very lowest trade intensities). Further evidence on the importance of state-dependency is provided by the four test statistics shown at the bottom of the panel. Here we report the results of Wald tests for the following coefficient restrictions: (i) $\underline{d}_j(0) = \bar{d}_j(0)$, (ii) $\underline{d}_j(1) = \bar{d}_j(1)$, (iii) $\bar{d}_j(1) = \bar{d}_j(0)$, and (iv) $\underline{d}_j(1) = \underline{d}_j(0)$ for $j = \{1, \dots, 6\}$. As the table shows, there is strong statistical evidence against all of these restrictions. These findings are consistent with the non-parametric evidence on state-dependence in hourly price change data reported in Evans and Lyons (2002b). They also illustrate how intraday data brings greater resolution to the study of price and order flow dynamics than is possible with models estimated at lower frequencies.

Parameter estimates from the order flow equation are reported in Panel B. Many of the coefficients are highly statistically significant, indicating that there is indeed a good deal of serial correlation in intraday order flow. The table also reports the estimate of $(1 - \sum_j c_j)^{-1}$ which measures the cumulative long-run effect of dispersed information on order flow. The estimate of 1.69 indicates that the cumulative effect of a dispersed information shock is approximately 70 percent greater than its initial impact.

Panel C of Table 3 reports the estimated parameters of the state-dependent error variances. The estimated values for $\underline{\omega}_\eta(A)$ imply that the standard deviation of the idiosyncratic shocks slowly falls from approximately 8 to 3 per cent as n varies from 2 to 200. Thus, the cross-sectional dispersion of transactions prices falls as trade intensity increases, as in Evans (2002), but we find no evidence that dispersion depends on the arrival of news. The estimates of $\bar{\omega}_\varepsilon(A)$ indicate how the volatility of common-knowledge shocks varies with trade

Table 3: GMM Estimates of the Intraday Model

A: Price Equation: $\Delta p_i = \sum_{j=1}^6 \{ \underline{d}_j(\bar{A}_i) e^{-n_i/100} + \bar{d}_j(\bar{A}_i) (1 - e^{-n_i/100}) \} x_{i+5-j} + \varepsilon_i$							
	$\underline{d}_1(\cdot)$	$\underline{d}_2(\cdot)$	$\underline{d}_3(\cdot)$	$\underline{d}_4(\cdot)$	$\underline{d}_5(\cdot)$	$\underline{d}_6(\cdot)$	$\sum_j \underline{d}_j(\cdot)$
$\hat{A} = 0$	0.029 (0.024)	0.025 (0.057)	0.028 (0.233)	-0.047 (0.052)	-0.113 (0.025)	-0.034 (0.033)	-0.113 (0.030)
$\hat{A} = 1$	-0.022 (0.045)	0.074 (0.046)	0.054 (0.042)	-0.131 (0.046)	0.002 (0.044)	-0.066 (0.043)	-0.089 (0.070)
	$\bar{d}_1(\cdot)$	$\bar{d}_2(\cdot)$	$\bar{d}_3(\cdot)$	$\bar{d}_4(\cdot)$	$\bar{d}_5(\cdot)$	$\bar{d}_6(\cdot)$	$\sum_j \bar{d}_j(\cdot)$
$\hat{A} = 0$	0.127 (0.106)	0.275 (0.210)	0.543 (0.716)	0.629 (0.186)	-0.220 (0.078)	-0.062 (0.101)	1.293 (0.106)
$\hat{A} = 1$	0.278 (0.153)	-0.018 (0.139)	0.256 (0.131)	0.858 (0.133)	-0.449 (0.107)	0.091 (0.114)	1.015 (0.209)
Wald Tests							
	$\underline{d}_j(0) = \bar{d}_j(0)$	$\underline{d}_j(1) = \bar{d}_j(1)$	$\bar{d}_j(1) = \bar{d}_j(0)$	$\underline{d}_j(1) = \underline{d}_j(0)$			
	216.083 (<0.001)	19.096 (0.004)	20.896 (0.002)	11.953 (0.063)			
B: Order Flow Equation: $x_i = \sum_{j=1}^{10} c_j x_{i-j} + \xi_{i-4}$							
	c_1	c_2	c_3	c_4	c_5	c_6	
	0.21 (0.014)	0.036 (0.013)	0.048 (0.012)	0.033 (0.012)	0.019 (0.011)	0.025 (0.011)	
	c_7	c_8	c_9	c_{10}	$(1 - \sum_{j=1}^{10} c_j)^{-1}$		
	0.015 (0.010)	0.017 (0.012)	-0.016 (0.010)	0.020 (0.008)	1.688 (0.070)		
C: Variance Parameters: $\Omega_j(n, A) = \underline{\omega}_j(A) e^{-n_i/100} + \bar{\omega}_j(A) (1 - e^{-n/100})$							
Shock Types							
	<u>Idiosyncratic</u>		<u>Common Knowledge</u>		<u>Dispersed Information</u>		
	$\underline{\omega}_\eta(\cdot)$	$\bar{\omega}_\eta(\cdot)$	$\underline{\omega}_\varepsilon(\cdot)$	$\bar{\omega}_\varepsilon(\cdot)$	$\underline{\omega}_\xi(\cdot)$	$\bar{\omega}_\xi(\cdot)$	
$A = 0$	0.002 (<0.001)	0.000	0.000	0.010 (<0.001)	0.000	0.032 (0.002)	
$A = 1$	0.002 (<0.001)	0.000	0.000	0.006 (0.002)	0.000	0.032 (0.034)	

Notes: The table reports GMM estimates with asymptotic standard errors in parentheses corrected for conditional heteroskedasticity and an MA(1) error term. News arrival is denoted by A_i and \bar{A}_i , with $\bar{A}_i = \max\{A_i, A_{i-1}, A_{i-2}\}$ where $A_i = 1$ if there was a news arrival during the previous 5-minutes. Coefficients and standard errors in panels A and C are multiplied by 100. P-values are reported in parentheses below the Wald statistics in panel A. For the variance parameters, P-values are not reported in cases where unrestricted parameter estimates were <0.0001 because these parameters were restricted to zero.

intensity and the arrival of news. The estimated standard deviation of common-knowledge shocks rises from approximately 1 to 9 percent as n varies between 2 and 200 when news is absent, and from 1 to 7 percent when news arrives. The estimated standard deviation of dispersed information shocks also increases with trade intensity: from 0.03 to 0.17 as n varies between 2 and 200, whether or not news arrives.

Two implications of these estimates deserve emphasis. First, under normal trading conditions, much of the observed volatility in high frequency transactions prices is attributable to the dispersion of prices that characterizes market activity at a point in time.¹³ Failure to account for this feature of the data would leave our analysis of how news arrivals affect prices and order flow flawed. Second, our estimates only show how the arrival of news affects price and order flow dynamics for a given level of trade intensity. If the arrival of news changes trade intensity, as indeed it does, the total impact of news on prices and order flow will reflect both the direct effect of news and the indirect effects associated with the induced change in trade intensity. We examine the combined effects of news in Table 5 below.

Table 4: Diagnostics for Intraday Model

	(a) $C(L)$	(b) $D(L, S)$	(c) $\Omega_\varepsilon(S)$	(d) $\Omega_\xi(S)$	(e) $\Omega_\eta(S)$
<u>Instrument: z_i</u>					
(i) Trade Intensity n_i	2.624 (0.989)	0.371 (0.999)	0.737 (0.391)	0.200 (0.655)	0.134 (0.714)
(ii) US News A_i^{US}	1.731 (0.188)	19.083 (0.087)	0.950 (0.330)	1.731 (0.188)	0.019 (0.891)
(iii) Scheduled News A_i^{S}	17.905 (0.084)	13.084 (0.363)	3.904 (0.068)	2.307 (0.129)	1.660 (0.198)
(iv) Residual ARCH			17.543 (0.001)	30.123 (<0.001)	8.190 (0.042)

Notes: The table reports C -tests for a set of auxiliary moment conditions implied by the model. In column (a) the restrictions take the form $E[\xi_i x_{i+4-j} z_i] = 0$ for $j = \{1, 2, \dots, 10\}$. The restrictions in (b) are $E[u_i^o z_i x_{i+5-j}] = 0$ for $j = \{1, 2, \dots, 6\}$ where $u_i^o \equiv \varepsilon_i + \eta_i^o - \eta_{i-1}^o$ for $o = \{\text{ask}, \text{bid}\}$. In columns (c) - (e) the restrictions are $E[\varkappa_i z_i] = 0$ where $\varkappa_i \equiv u_i^o u_i^o - \Omega_\varepsilon(n_i, A_i)$ in (c), $\varkappa_i \equiv \xi_i^2 - \Omega_\xi(n_i, A_i)$ in (d), and $\varkappa_i \equiv u_i^o u_{i-1}^o + \Omega_\eta(n_{i-1}, A_{i-1})$ in (e). The instruments z_i are n_i , A_i^{US} , A_i^{S} and \varkappa_{i-j} for $j = \{1, 2, 3\}$ in rows (i) - (iv) respectively. P-values are reported in parentheses.

Our specification for the intraday model imposes many more moment conditions than were used in GMM estimation. Table 4 provides diagnostics in the form of C-tests on a selection of these additional moment conditions. The tests in column (a) look for state-dependency in the order flow polynomial $C(L)$. For this

¹³Evans (2002) describes how the low degree of transparency of the FX institutional structure allows for the existence of a price distribution without introducing arbitrage opportunities.

purpose we compute C -statistics for restrictions of the form $E[\xi_i x_{i+4-j} z_i] = 0$ for $j = \{1, 2, \dots, 10\}$; where z_i equals n_i , A_i^{US} and A_i^{S} in rows (i) (ii), and (iii) respectively. These moment conditions will not hold if, contrary to the assumption of our model, the serial correlation in order flow varies with either trade intensity, the arrival of US news, or the arrival of scheduled news. The tests reported in column (b) look for misspecification in the estimated form of the $D(L, S)$ polynomial. In this case the restrictions being tested take the form $E[u_i^o z_i x_{i+5-j}] = 0$ for $j = \{1, 2, \dots, 6\}$ where $u_i^o \equiv \varepsilon_i + \eta_i^o - \eta_{i-1}^o$ for $o = \{\text{ASK, BID}\}$. These tests look for evidence of state-dependency in $D(L, S)$ beyond that implied by functional form in (7) and (8). Similarly, the C -tests in columns (c)-(e) look for evidence of misspecification in the error variances. The moments being tested here take the form of $E[\varkappa_i z_i] = 0$ where \varkappa_i is the unexpected squared realization of the shock in period i [i.e., $\varkappa_i \equiv u_i^o u_i^o - \Omega_\varepsilon(n_i, A_i)$ in column (c), $\varkappa_i \equiv \xi_i^2 - \Omega_\xi(n_i, A_i)$ in (d), and $\varkappa_i \equiv u_i^o u_{i-1}^o + \Omega_\eta(n_{i-1}, A_{i-1})$ in (e)]. Row (iv) reports C -tests for 3rd order residual ARCH by testing moment conditions of the form $E[\varkappa_i \varkappa_{i-j}] = 0$ for $j = \{1, 2, 3\}$.

As the table shows, none of the test statistics in rows (i)-(iii) are significant at the 5 percent level. In particular, there is no evidence from the tests in row (i) that the functional forms in (7)-(9) are unduly restrictive. The results in rows (ii) and (iii) address the question of whether there should be a distinction in our model between the arrival of US and German news, or scheduled and unscheduled news. Recall that the median (daily) arrival rate for German news is four times the rate for US news. Some of this difference may be attributable to institutional features, such as the distribution of news bureaus supplying Reuters, that are unrelated to the pace at which price-relevant information becomes known. In particular, it is possible that the arrival rate of German news items on the Headline screens overstates the true pace at which price-relevant German news arrives. In this case, our specification using the A_i dummy will overstate how the dynamics of prices and order flow change immediately following the arrival of price-relevant news. The C -statistics in row (ii) test for this form of misspecification using the arrival of US news as an instrument. None of the statistics are significant at the 5 percent level. Differences between the arrival of scheduled and unscheduled news could pose similar problems. For example, if the ratio of common-knowledge to dispersed information in scheduled news is higher on average than in non-scheduled news, the price and order flow dynamics following the arrival of scheduled news may differ from the dynamics following the arrival of other news. The C -statistics in row (iii) are designed to look for evidence of this form of misspecification. None are significant at the 5 percent level.¹⁴ In sum, these diagnostic tests suggest that the estimated model adequately accounts for the effects of varying trade intensity and the arrival of news on the dynamics of transaction prices and interdealer order flow.

The model is less successful in accounting for all the heteroskedasticity in the error processes. The C -tests for 3rd-order residual ARCH are significant at the 5 percent level. An inspection of the estimated residuals

¹⁴ Since the arrival of scheduled news is, by definition, exogenous to past market volatility, these results are consistent with the absence of feedback from FX price volatility to the arrival of unscheduled news items. We also looked more directly for evidence of feedback by estimating logit and probit models for A_i and A_i^{US} and A_i^{GM} using lagged square price changes, specifically $\{\Delta p_{i-j}^{\text{ASK}}\}_{j=6}^{24}$, as explanatory variables. In all cases, the estimated coefficients were small and statistically insignificant. There is no evidence of feedback effects in our filtered series of unscheduled news items.

shows that these residual ARCH effects are concentrated at lag one. In fact, if we omit this moment from our C -test, we cannot reject the null of no residual heteroskedasticity. We have accounted for this feature of the data in our estimates and tests by constructing the GMM weighting matrix from the Newey West estimator with an MA(1) serial correlation correction.¹⁵

3.1.4 News Arrival and Intraday Dynamics

We now examine how the information in macro news is transmitted to prices. For this, we use our model estimates to compute a variance decomposition for price changes across different market states. First, we use our estimates to write the change in average transaction price as:

$$\Delta p_i = B(L, S_i)\xi_i + \varepsilon_i, \quad (14)$$

where $B(L, S) = D(L, S)C(L)L^m$. The state-dependent coefficients in $B(L, S)$ identify how dispersed information affects prices and can be computed from our estimates of the coefficients in $D(L, S)$ and $C(L)$. We can also use equation (14) to decompose the variance of price changes into different theoretical components. In particular, consider the k -period price change between period $i - k$ and i : $\Delta^k p_i \equiv \sum_{j=0}^{k-1} \Delta p_{i-j}$. Substituting for Δp_i with (14), gives:

$$\Delta^k p_i = \sum_{j=0}^{k-1} \varepsilon_{i-j} + \sum_{j=0}^{k-1} B(L, S_{i-j})\xi_{i-j}. \quad (15)$$

Since the ξ_i , and ε_i shocks are mutually independent and serially uncorrelated, (15) implies that:

$$Var\left(\Delta^k p_i \mid \{S_{i-j}\}_{j=0}^{k-1}\right) = \sum_{j=0}^{k-1} \Omega_\varepsilon(S_{i-j}) + \sum_{j=0}^{k-1} B(L, S_{i-j})^2 \Omega_\xi(S_{i-j}). \quad (16)$$

Equation (16) provides a decomposition of the variance of price changes conditioned on the state of the market during the last k periods. The first component on the right-hand side is the variance contribution of common-knowledge shocks, the second is the contribution of dispersed information shocks operating via order flow. Notice that state-dependency in the error variances and lag polynomial $D(L, S)$ of our model allows the contribution of each variance component to vary with changes in trade intensity and the arrival of macro news. We now use the model estimates to quantify these effects.

Order flow is much more important in price determination when macro news arrives. Table 5 reports the estimated contribution of dispersed information to the variance of price changes over horizons of 5, 30 and 60 minutes (i.e., $k = \{1, 6, 12\}$) when trading intensity is at four different levels (i.e., $n = \{5, 10, 20, 30\}$). Row (i) in each panel reports the contribution for a given level of trade intensity in the absence of macro news. (The statistics in parenthesis are standard errors associated with these estimates computed from the asymptotic distribution of the GMM estimates by the ‘‘delta-method’’.¹⁶) Consistent with the results in Evans (2002),

¹⁵Specifically, the presence of first-order ARCH induces serial correlation in the residuals associated with conditions (10b), (11b), (11c) and (11d).

¹⁶Specifically, let $R^k(\theta, n, A)$ denote the contribution of dispersed information shocks equal to

these statistics show that the contribution of dispersed information to price variance rises with trade intensity and horizon. The contribution of dispersed information in the presence of macro news is reported in row (ii). These statistics incorporate direct effects of news arrival via the A and \bar{A} dummies and the indirect effects via the induced change in trade intensity. We estimate that trading intensity rises by approximately 9 trades per minute when news arrives.¹⁷ To estimate the contribution of dispersed information we therefore use the GMM estimates of (16) with $B(L, S^A)$, $\Omega_\xi(S^A)$, and $\Omega_\varepsilon(S^A)$ where $S^A = \{n + 9, 1\}$ and n is the initial level of trade intensity shown at the top of each panel in the table. A comparison of the statistics in rows (i) and (ii) show that following the arrival of macro news, dispersed information contributes more to the variance of prices across all three horizons. This pattern also appears consistently across all four panels (corresponding to different initial levels of trade intensity).

We conducted a Monte Carlo experiment to assess the statistical significance of these findings. The experiment comprised the following steps: (i) draw a vector of parameter estimates $\hat{\theta}^j$ from the estimated asymptotic distribution of the GMM estimates; $N(\hat{\theta}, \hat{V}_\theta)$, (ii) use (16) and $\hat{\theta}^j$ to compute the contribution of the dispersed information shocks to the k -period price variance at trade intensity n in the absence of news ($A = \bar{A} = 0$), $R^k(\hat{\theta}^j, n, 0)$ for horizons of 5, 30 and 60 minutes (i.e., $k = \{1, 6, 12\}$), (iii) use (16) and $\hat{\theta}^j$ to compute the contribution to k -period price variance with news ($A = \bar{A} = 1$) at trade intensity $n^A = n + 9$, $R^k(\hat{\theta}^j, n^A, 1)$ for $k = \{1, 6, 12\}$, and (iv) repeat steps (i) - (iii) 5000 times for $n = \{5, 10, 20, 30\}$ and compute the fraction of times that $R^k(\hat{\theta}^j, n, 0) \geq R^k(\hat{\theta}^j, n^A, 1)$. This procedure gives us a Monte Carlo estimate of the p-value for the null hypothesis that news arrival does not increase the contribution of dispersed news to the variance of prices. Cases where the p-values are less than 10, 5 and 1 percent are indicated in Table 5 by “**”, “***”, and “****” respectively. Based on these calculations, the increased contribution of dispersed information shocks following the arrival of macro news is strongly significant over most horizons and initial trading intensities.

The specification tests reported in Table 4 do not suggest that the direct affects of macro news arrival vary according to whether or not the news item is a scheduled. Nevertheless, scheduled news may have a different *total* impact because the induced trade intensity differs from the trade intensity induced by non-scheduled news. We estimate that trading intensity when scheduled U.S. news arrives rises by approximately 13 trades per minute. Row (iii) of Table 5 shows the contribution of dispersed information in the presence of a scheduled news announcement that increases trade intensity by this amount. Because the price-impact of order flow increase with trading intensity, the estimated variance contribution of dispersed information is larger following the arrival of scheduled news than it is for the more prevalent non-scheduled items. The

$\left\{ \sum_{j=0}^{k-1} B(L, S)^2 \Omega_\xi(S) \right\} \left\{ \sum_{j=0}^{k-1} \Omega_\varepsilon(S) + \sum_{j=0}^{k-1} B(L, S)^2 \Omega_\xi(S) \right\}^{-1}$ given a constant level of trading intensity n , and the presence or absence of macro news, $A = \{1, 0\}$. We estimate the standard error of $R^k(\theta, n, A)$ as the square root of $\nabla R^k(\hat{\theta}, n, A)' \hat{V} \nabla R^k(\hat{\theta}, n, A)$ where $\nabla R^k(\cdot)$ is the gradient vector w.r.t. θ , and \hat{V} is the estimated covariance matrix of the GMM estimates, $\hat{\theta}$.

¹⁷This estimate is obtained from the OLS estimate of δ from the regression: $n_i = \delta A_i + \sum \gamma_i dum_{i,\tau} + u_i$ where $dum_{i,\tau}$ is a “seasonal” time dummy that takes the value of one when observation i falls in the τ ’th 30-minute window of a day. We estimate δ to be 8.91 with a standard error of 0.62.

Table 5: Variance Decomposition

	Horizon (minutes)			Horizon (minutes)		
	5	30	60	5	30	60
	trade intensity: $n = 5$			trade intensity: $n = 10$		
(i) No News	0.631 (1.040)	0.989 (2.811)	0.758 (3.754)	1.436 (1.327)	2.314 (2.911)	2.118 (3.621)
(ii) News	3.895** (0.911)	10.280** (3.396)	11.768* (4.236)	5.123** (1.354)	12.137** (4.554)	13.597** (5.451)
(iii) Scheduled News	8.271*** (2.896)	16.083** (8.020)	17.417* (9.112)	9.868*** (3.748)	17.807** (9.569)	19.067** (10.727)
	trade intensity: $n = 20$			trade intensity: $n = 30$		
(i) No News	3.808 (1.359)	7.475 (2.850)	7.957 (3.303)	7.173 (1.738)	14.862 (3.747)	16.129 (4.131)
(ii) News	7.981** (2.755)	15.754* (7.658)	17.101* (8.729)	11.214** (4.500)	19.163 (10.673)	20.358 (11.862)
(iii) Scheduled News	13.231*** (5.533)	21.067* (12.326)	22.163* (13.573)	16.679** (7.248)	24.053 (14.569)	24.980 (15.871)

Notes: The table reports values for $R^k(\theta, n, A)$, the contribution of dispersed information shocks to variance of k -horizon price changes implied by the GMM estimates of the intraday model given a constant level of trading intensity n , and the presence or absence of macro news, $A = \{1, 0\}$. Standard errors are in parentheses. Statistics in rows (i) - (iii) are computed as $R^k(\theta, n, 0)$, $R^k(\theta, n + 9, 1)$ and $R^k(\theta, n + 13, 1)$ respectively. Cases where the Monte Carlo p-value for the null that news arrival does not increase the contribution of dispersed news to the variance of prices is less than 10, 5 and 1 percent are indicated by “*”, “**”, and “***” respectively.

p-values computed from Monte Carlo experiments with $n^A = n + 13$ indicate an even stronger pattern of statistical significance.

Overall, our estimates indicate that order flow contributes more to price adjustment following macro news than at other times. This is not what one would expect if macro news is primarily comprised of common-knowledge information that is directly impounded into FX prices. If macro news primarily transmits new common-knowledge information, order flow should contribute less to price-dynamics in the period following the arrival of news than at other times. By contrast, the results in Table 5 strongly suggest that the arrival of macro news triggers trading that reveals new dispersed information that affects prices indirectly. One particularly interesting aspect of our findings concerns the effects of scheduled U.S. announcements. Since these news items contain data releases on macro economic aggregates, one might have expected that they contain a greater proportion of common-knowledge to dispersed information than some of the other news items in our sample. That order flow is at least as important in price dynamics following scheduled news suggests that this common view concerning the information content of macro news is incorrect.

4 Daily Analysis

Our intraday analysis shows the importance of the order flow channel as a means for impounding macro news in FX prices. We now examine implications of this for the behavior of FX prices at the daily frequency. This examination compliments our intraday analysis for three reasons. First, daily changes in FX prices are very nearly a martingale (which is not true of five-minute changes). Our daily model thus sheds light on how the information contained in macro news contributes to price variation over the longer run. Second, our daily analysis provides additional perspective on results relating daily price dynamics to order flow (e.g., Evans and Lyons 2002a). In particular, our estimates provide a breakdown of the sources of price and order flow volatility. Third, our daily analysis provides a robustness check on the results presented above. For example, we can construct measures of the daily flow of macro news in ways that were not possible at higher frequencies. The consistency of the results derived from estimates of the daily and intraday model shows that our main findings are robust to our methods for identifying the impact of macro news arrivals.

4.1 The Model

Our daily model for price and order flow dynamics comprises the following equations:

$$\Delta p_t = \alpha x_t + e_t + v_t, \tag{17}$$

$$x_t = u_t + w_t, \tag{18}$$

where Δp_t is the change in the spot price of FX between 5:00 pm on day $t - 1$ and 5:00 pm on day t and x_t is interdealer order flow realized over the same period. The parameter α captures the price impact of order flow at the daily horizon, i.e., it reflects information content. Prices and order flow are subject to four shocks representing different sources of information hitting the market: e_t, v_t, u_t , and w_t . These shocks are mean zero, mutually uncorrelated, and serially uncorrelated. The e_t and v_t shocks represent information that is impounded in price directly. e_t is the common knowledge effect of macro news arrivals on the price of FX. v_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 u_t and w_t shocks. The u_t shocks represent order flow effects from macro news arrivals – the dispersed information effect of the news. Shocks to order flow that are unrelated to macro news are represented by the w_t shocks (e.g., portfolio shifts arising from other sources such as changing risk tolerances or hedging).

We identify the effects of the news-related common-knowledge and dispersed-information shocks, e_t and u_t , through state-dependency of price changes and order flow in the second moments. Specifically, we assume that the variance of e_t and u_t on day t is increasing in the daily flow of macro news, which we measure by the number of US and German news arrivals between 5:00 pm on days $t - 1$ and t , A_t^{US} and A_t^{G} :

$$\text{Var}_t(e_t) = \Sigma_e^2(A_t^{\text{US}}, A_t^{\text{G}}), \text{ and } \text{Var}_t(u_t) = \Sigma_u^2(A_t^{\text{US}}, A_t^{\text{G}}), \tag{19}$$

where $\Sigma_{\varkappa}^2(0,0) = 0$, with $\partial\Sigma_{\varkappa}^2/\partial A_t^k > 0$ for $\varkappa = \{e, u\}$ and $k = \{\text{US}, \text{G}\}$. Thus, on days without news, $e_t = u_t = 0$, so price changes and order flow are driven solely by the v_t and w_t shocks. These shocks are independent of news, so their variances are unrelated to A_t^k . As we shall see, there is little evidence of state-dependency in the second moments of daily price changes and order flow beyond the effects of news. In particular, unlike our intraday model, there is no need to incorporate trade intensity as an additional state variable. We therefore assume that the conditional variances of the v_t and w_t shocks are constant:

$$\text{Var}_t(v_t) = \Sigma_v^2, \text{ and } \text{Var}_t(w_t) = \Sigma_w^2. \quad (20)$$

Several features of our daily model deserve comment. First, our specification abstracts from the complex intraday dynamics of prices and order flow. Equations in (17) and (18) imply that by 5:00 pm GMT each day, FX prices fully reflect the information contained in order flow to that point. As a result, price change over the next 24 hours (i.e. Δp_{t+1}) are not correlated with order flow from the past 24 hours, (i.e. x_t). This feature of our model is supported by the data. We show below that there is no correlation between Δp_{t+1} and x_t . Our specification also implies the absence of serial correlation in daily price changes and order flows. This too is consistent with the evidence reported in Section 2. A second feature of our specification concerns the price-impact parameter α . Our intraday analysis showed that the price impact of order flows varied with trade intensity and the arrival of news. This form of state-dependency in the intraday data does not appear at the daily frequency (addressed below), so we do not allow for state-dependency in α . We would add that this restriction in our model means that our test of the relative importance of indirect effects is conservative: order flow induced by news may have more price impact than the constrained equation gives it credit for. In any event, we do incorporate state-dependency into the error variances. This final feature is key to identifying the effects of macro news, so let us focus on it more closely.

Identification of the effects of macro news is achieved by the assumption that the variance of the e_t and u_t shocks is higher on days when there are a greater number of news items appearing on the Reuters Money Market News screen. Crucially, this assumption does not require that FX market participants view the information in each news item as equally important (which the market does not). The identifying power of this assumption does, however, depend on the absence of wild variations in the quality of Reuters' editorial judgements. For example, if the Reuters screen were flooded one day with reports containing essentially no information, but on another a few reports appeared with great economic significance, daily variations in the number of news reports would be a poor measure of the daily flow of macro news. Based on our understanding of Reuters' editorial process, this possibility seems far-fetched. That said, we recognize that no single measure will identify the daily variation in macro news flow with complete precision. Thus, in addition to measures based on the daily arrival rates for US and German news shown in (19), we will also use measures based on the subset of items that are scheduled.

4.2 Estimation

We estimate two versions of the model by the Generalized Method of Moments. Version I assumes that the variance of the e_t and u_t shocks on day t varies only with the sum of the US and German news items, $A_t^{\text{ALL}} \equiv A_t^{\text{US}} + A_t^{\text{G}}$. Under this specification, the flow of macro news is identified by the arrival rate of both US and German news. We also allow for the possibility that daily variations in the flow of macro news may be reflected differently in the arrival rates for US and German news. Version II of our model allows the variance of e_t and u_t on day t to depend on the number of US and German news items separately. The variance functions are assumed to be linear in both versions of the model:

$$\text{Version I: } \Sigma_{\varkappa}^2(A_t^{\text{US}}, A_t^{\text{G}}) = \sigma_{\varkappa} A_t^{\text{ALL}} \quad (21)$$

$$\text{Version II: } \Sigma_{\varkappa}^2(A_t^{\text{US}}, A_t^{\text{G}}) = \sigma_{\varkappa}^{\text{US}} A_t^{\text{US}} + \sigma_{\varkappa}^{\text{G}} A_t^{\text{G}}$$

where $\sigma_{\varkappa}, \sigma_{\varkappa}^{\text{US}}$ and $\sigma_{\varkappa}^{\text{G}}$ are positive parameters for $\varkappa = \{e, u\}$. Thus, the parameters to be estimated are $\{\alpha, \Sigma_w^2, \Sigma_v^2, \sigma_e, \sigma_u\}$ in Version I, and $\{\alpha, \Sigma_w^2, \Sigma_v^2, \sigma_e^{\text{US}}, \sigma_e^{\text{G}}, \sigma_u^{\text{US}}, \sigma_u^{\text{G}}\}$ in Version II.

The GMM estimates of the model parameter are derived from the following set of moment conditions:

$$0 = E[(\Delta p_t - \alpha x_t) x_t] \quad (22a)$$

$$0 = E[\{\mathcal{V}_t(\Delta p_t) - \text{Var}_t(\Delta p_t)\} \otimes \mathcal{Z}_t], \quad (22b)$$

$$0 = E[\{\mathcal{V}_t(x_t) - \text{Var}_t(x_t)\} \otimes \mathcal{Z}_t], \quad (22c)$$

where \mathcal{Z}_t is a vector of instruments. Condition (22a) follows from the assumed orthogonality between the shocks to prices (e_t and v_t) and the shocks to order flow (u_t and w_t). Conditions (22b) and (22c) combine the second moments of price changes and order flow implied by the model with measures of the variance of order flow, $\mathcal{V}(x_t)$, the variance of price changes, $\mathcal{V}(\Delta p_t)$. These measures are computed for each day in our sample from the 5-minute intraday observations as:

$$\mathcal{V}_t(\Delta p_t) = \sum_{i=1}^{288} \Delta p_{it}^2, \quad \mathcal{V}_t(x_t) = \sum_{i=1}^{288} x_{it}^2, \quad (23)$$

where the subscript “ it ” denotes the i 'th 5-minute observation on day t . $\mathcal{V}_t(\Delta p_t)$ and $\mathcal{V}_t(x_t)$ are the (un-centered) second moments of the price change and order flow process over day t , scaled by the number of 5-minute intraday observations. Andersen, Bollerslev, Diebold, and Labys (2001) show that these measures are consistent nonparametric estimates of the actual moments under mild regularity conditions. They also note that while the measures will be biased when prices changes and order flow do not follow Martingales in the continuous time limit, in practice these biases will be very small if a large number of high frequency observations are used to compute each daily measure. This appears true in our data. Estimates of $\mathcal{V}_t(\Delta p_t)$, and $\mathcal{V}_t(x_t)$ computed from Δp_{it} and x_{it} are almost identical to their counterparts using the estimated resid-

uals from the price and order flow equations of the intraday model: the correlation between the alternative measures is greater than 0.99 for both order flow and price changes.¹⁸

We use two sets of instruments to implement estimation. The instrument vector in Version I comprises a constant and sum of the US and German news items, A_t^{ALL} . In Version II, we use a constant, A_t^{US} and A_t^{G} as instruments. These choices imply that the number of moment conditions in (22a-c) equals the number of parameters, so the estimates come from exactly identified versions of model. As above, we apply the standard 2-step method to compute the GMM estimates (without the serial correlation correction in the weighting matrix). We will also consider the adequacy of our model estimates with a set of diagnostic tests based on additional moment conditions.

4.2.1 Daily Estimates

Panel A of Table 6 reports parameter estimates from both versions of the model with exact identification. Asymptotic standard errors allowing for residual heteroskedasticity are shown in parentheses. 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.) In Version I of the model, both variance parameters σ_e and σ_u are positive and significant at the five percent level. These estimates imply that both direct and indirect effects of news on price are present. This finding is confirmed by the estimates from Version II reported in the right-hand panel. When U.S. and German news events are introduced separately, the estimates of σ_e^{US} , σ_e^{G} , σ_u^{US} , and σ_u^{G} are all positive and significant at the five percent level. Furthermore, as panel B shows, Wald statistics for the null that $\sigma_e^{\text{US}} = \sigma_u^{\text{US}} = 0$, and $\sigma_u^{\text{US}} = \sigma_u^{\text{G}} = 0$, are highly significant. Panel B also shows that there is no significant evidence against the parameter restrictions imposed by Version I of the model, namely $\sigma_e^{\text{US}} = \sigma_u^{\text{US}}$ and $\sigma_u^{\text{US}} = \sigma_u^{\text{G}}$.

Panel C shows results of diagnostic tests that examine an expanded set of moment conditions provide additional support for our specification. In row (i) we report the J -statistic for specifications using (22) and $E[(\Delta p_t - \alpha x_t) x_{t-1}] = 0$ as moment conditions.¹⁹ Our model should satisfy this additional condition because all the price impact of order flow occurs within the day. As the table shows, there is no significant evidence to reject this set of restrictions in either version of the model. The statistics in row (ii) test for the presence of (residual) serial correlation in the price change and order flow process by respectively adding $E[(\Delta p_t - \alpha x_t)(\Delta p_{t-1} - \alpha x_{t-1})] = 0$ and $E[x_t x_{t-1}] = 0$ to the conditions in (22). Again, consistent with the assumed structure of our model, none of the J -statistics are statistically significant. Next, we turn to the issue of state-dependency. Our daily model assumes that trade intensity and news have no effect on α , the parameter identifying the price-impact of order flow. We examine this restriction by adding $E[(\Delta p_t - \alpha x_t) \otimes z_t] = 0$ to the conditions in (22) for $z_t = \{x_t n_t, x_t A_t^{\text{ALL}}\}$ in Version I and

¹⁸This finding is not inconsistent with the results of intraday model. Although intraday order flow is serially correlated, lagged order flow accounts for a very small fraction of the variation in realized order flow over a 5-minute period. Similarly, the leads and lags in intraday order flow account for a small fraction of the variance in high frequency price changes.

¹⁹The J -statistics reported here are equivalent to the C-statistics used in our intraday analysis because both versions of the daily model are exactly identified without the additional moment conditions.

Table 6: GMM Estimates of Daily Models

A: Parameters	<u>Version I</u>		<u>Version II</u>	
	Estimate	Std. Err	Estimate	Std. Err
α	0.032	(0.003)	0.032	(0.003)
σ_w^2	67.231	(11.395)	67.018	(11.282)
σ_v^2	3.530	(0.675)	3.518	(0.671)
σ_e	3.737	(0.813)		
σ_u	0.188	(0.053)		
σ_e^{US}			5.682	(2.661)
σ_e^{G}			3.358	(0.977)
σ_u^{US}			0.291	(0.147)
σ_u^{G}			0.168	(0.063)

B: Wald Tests	Statistic	p-value
$\sigma_e^{\text{US}} = \sigma_u^{\text{US}} = 0$	33.303	(0.000)
$\sigma_u^{\text{US}} = \sigma_u^{\text{G}} = 0$	13.707	(0.001)
$\sigma_e^{\text{US}} = \sigma_u^{\text{US}} \ \& \ \sigma_u^{\text{US}} = \sigma_u^{\text{G}}$	0.763	(0.683)

C: Diagnostic Tests	Statistic	p-value	Statistic	p-value
i) Lagged order flow	2.502	(0.114)	2.502	(0.114)
ii) Serial correlation:				
Δp_t eqn.	0.014	(0.905)	0.014	(0.905)
x_t eqn.	0.190	(0.663)	0.190	(0.663)
iii) State-dependency:				
α	2.767	(0.251)	2.767	(0.251)
$\text{Var}(\Delta p_t) \ \& \ \text{Var}(x_t)$	2.479	(0.290)	2.527	(0.283)
iv) Residual Arch:				
Δp_t eqn.	0.348	(0.555)	0.281	(0.596)
x_t eqn.	2.332	(0.127)	2.486	(0.115)
v) Joint Test	10.097	(0.343)	9.876	(0.361)

Notes: Panel A of the table reports GMM parameter estimates and asymptotic standard errors (corrected for heteroskedasticity) in parentheses. Panel B shows Wald tests for the coefficient restrictions listed on the left with asymptotic p-values reported in parentheses. The J -tests shown in panel C test the moment restrictions in (22) and the following: (i) $E[(\Delta p_t - \alpha x_t) x_{t-1}] = 0$, (ii) $E[(\Delta p_t - \alpha x_t)(\Delta p_{t-1} - \alpha x_{t-1})] = 0$, $E[x_t x_{t-1}] = 0$, (iii) $E[(\Delta p_t - \alpha x_t) \otimes z_t] = 0$, $E[\{\mathcal{V}_t(\Delta p_t) - \text{Var}_t(\Delta p_t)\} n_t] = 0$, and $E[\{\mathcal{V}_t(x_t) - \text{Var}_t(x_t)\} n_t] = 0$, where $z_t = \{x_t n_t, x_t A_t^{\text{ALL}}\}$ in Version I and $z_t = \{x_t n_t, x_t A_t^{\text{US}}, x_t A_t^{\text{G}}\}$ in Version II, (iv) $E[\{\mathcal{V}_t(\Delta p_t) - \text{Var}_t(\Delta p_t)\} \{V_{t-1}(\Delta p_{t-1}) - \text{Var}_{t-1}(\Delta p_{t-1})\}] = 0$ and $E[\{\mathcal{V}_t(x_t) - \text{Var}_t(x_t)\} \{V_{t-1}(x_{t-1}) - \text{Var}_{t-1}(x_{t-1})\}] = 0$, and (v) all the moments listed in (i) - (iv). Asymptotic p-values are reported in parentheses.

$z_t = \{x_t n_t, x_t A_t^{\text{US}}, x_t A_t^{\text{G}}\}$ in Version II, where n_t denotes trading intensity on day t . As the table shows, neither of the associated J -statistics are significant. We also check for additional state-dependency in the error variances. In this case we add $E[\{\mathcal{V}_t(\Delta p_t) - \text{Var}_t(\Delta p_t)\} n_t] = 0$ and $E[\{\mathcal{V}_t(x_t) - \text{Var}_t(x_t)\} n_t] = 0$ to the conditions in (22). These additional moments examine whether the residual variance in price and order flow, unaccounted for by the arrival of news, is correlated with daily trade intensity. Once again, neither of the J -statistics is significant. There is no evidence that trade intensity should be present as a second state variable governing the error variances. Further evidence on the specification of the error variances is provided by the statistics in row (iv). Here we test for residual first order ARCH by adding $E[\{\mathcal{V}_t(\Delta p_t) - \text{Var}_t(\Delta p_t)\} \{\mathcal{V}_{t-1}(\Delta p_{t-1}) - \text{Var}_{t-1}(\Delta p_{t-1})\}] = 0$ and $E[\{\mathcal{V}_t(x_t) - \text{Var}_t(x_t)\} \{\mathcal{V}_{t-1}(x_{t-1}) - \text{Var}_{t-1}(x_{t-1})\}] = 0$ to the conditions in (22). These specification tests also show no evidence of significant misspecification in the error variances.²⁰ Finally, in row (v), we report J -statistics for models using (22) and all the additional moments. These moment conditions respectively provide 9 and 11 over-identifying restrictions in Versions I and II of the model. As the table shows, neither J -statistic is significant at the 5 percent level. The parameter estimates obtained in this manner are very similar to those reported in Panel A. Since the estimated standard errors are a little smaller (as one would expect), the overall pattern of statistical significance we report appears robust to the number of over-identifying restrictions used in estimation. Importantly this level of robustness is also reflected in the model-based statistics we consider next.

4.3 News Arrival and Daily Dynamics

Our intraday analysis showed that dispersed information contributes more to the variance of price changes following macro news announcements than at other times. Our daily model allows us to address a distinct but equally important issue: the extent to which macro news is impounded in prices directly, via the common knowledge e_t shocks, or indirectly via the dispersed information u_t shocks that affect prices via order flow.

To clarify this issue within the context of our daily model, consider the unconditional variance of price changes implied by our model, $\text{Var}(\Delta p_t)$. By definition, this variance can be written as $E[\text{Var}_t(\Delta p_t)] + \text{Var}(E_t \Delta p_t)$ where $E_t \Delta p_t$ and $\text{Var}_t(\Delta p_t)$ denote the first and second moments of price changes conditioned on the day t state of the market. According to our model, the number of news arrivals has no implication for the direction of how prices will change, so $E_t \Delta p_t = 0$. With the aid of equation (17), we can therefore write the unconditional variance as:

$$\text{Var}(\Delta p_t) = \alpha^2 E[\text{Var}_t(x_t)] + E[\text{Var}_t(e_t + v_t)].$$

²⁰An earlier version of this paper examined two further aspects of the model. We looked for evidence of nonlinearity in the error-variance specifications shown in (21) by regressing $\mathcal{V}_t(\Delta p_t)$ and $\mathcal{V}_t(x_t)$ on a constant, A_t^{US} , A_t^{G} , $(A_t^{\text{US}})^2$, and $(A_t^{\text{G}})^2$. Since the price and order flow variances are linear functions of the error variances, nonlinearity in the latter should appear in the form of non-zero coefficients on $(A_t^{\text{US}})^2$ and $(A_t^{\text{G}})^2$ in these regressions. Our estimates of these coefficients were not statistically significant. We also explored whether temporal aggregation could affect our results by introducing a feedback from price changes to order flow. Model estimates incorporating this feedback effect were similar to those reported here, and had the same implications concerning the effects of macro news.

The first term on the right identifies the contribution of order flow volatility to the variance of price changes. The second term identifies the contribution of information that is directly impounded into prices. Using equations (18)-(20) to substitute for $Var_t(x_t)$ and $Var_t(e_t + v_t)$, we obtain:

$$Var(\Delta p_t) = E[\Sigma_e^2(A_t^{US}, A_t^G)] + \alpha^2 E[\Sigma_u^2(A_t^{US}, A_t^G)] + \Sigma_v^2 + \alpha^2 \Sigma_w^2. \quad (24)$$

Equation (24) decomposes the unconditional variance of daily price changes into four components. The first term identifies the contribution of common-knowledge shocks associated with the arrival of news. We refer to this as the direct channel. The second term represents the contribution of dispersed information shocks associated with news. Notice that this term includes the price-impact coefficient α , because dispersed information affects prices via order flow. We refer to this as the indirect channel. The third and fourth terms identify the contribution of shocks that are not associated with the arrival of news; information embedded in the v_t and w_t shocks affects price via the direct and indirect channels respectively.

Table 7 reports elements of the variance decomposition in (24) derived from the estimates of the daily model. For this purpose, the expectations terms in (24) are replaced by sample averages (i.e., $E[\Sigma_{\varkappa}^2(A_t^{US}, A_t^G)]$ is replaced by $\frac{1}{T} \sum_{t=1}^T \hat{\Sigma}_{\varkappa}^2(A_t^{US}, A_t^G)$ for $\varkappa = \{e, u\}$). We also report standard errors computed by the “delta-method” from the estimated asymptotic distribution of the model estimates. The statistics shown in Panel A use the parameters estimated from the exactly identified models reported in panel A of Table 6. As noted above, these statistics are very similar to those based on the estimates derived from Versions I and II of the model with 9 and 11 over-identifying restrictions.

The upper rows in panel A of Table 7 report the contribution of dispersed and common knowledge information shocks to the unconditional variance of prices. The statistics in row (i) report the fraction of the unconditional variance attributable to the common knowledge shocks associated with news: $E[\Sigma_e^2(A_t^{US}, A_t^G)]/Var(\Delta p_t)$. Estimates from both versions of the model indicate that the direct effect of news arrivals account for approximately 14 percent of the variance of total price changes. The estimates from Version II of the model indicate that this total is split roughly 2 to 1 between German and US news. Since German news arrives at four times the daily rate of US news on average, these estimates suggest that a typical US news item has a somewhat larger direct effect on prices than a German item. Row (ii) reports the contribution of dispersed information to the variance of prices: $\alpha^2 E[\Sigma_u^2(A_t^{US}, A_t^G)]/Var(\Delta p_t)$. These statistics show that the indirect effects of news arrival account for roughly 22 percent of the variance. Once again, the arrival of German news contributes more than twice as much as US news through this channel. Row (iii) shows the total contribution of news to the variance of prices via both channels is approximately 36 percent. These estimates are an order of magnitude larger than those found in event studies. As noted in the introduction, the results from these studies imply that scheduled announcements account for only 2 to 3 per cent of the variance in daily price changes. The two main differences in approach are that here we consider both direct and indirect channels and we use a larger set of news items. Row (iv) shows the relative importance of the latter. Here we report the ratio of indirect to direct effects of news arrival implied by our

Table 7: Daily Price Variance Decompositions

A: All News	Version I	Version II		
	Combined	US	German	Combined
(i) Direct	0.139 (0.046)	0.036 (0.042)	0.104 (0.017)	0.140 (0.046)
(ii) Indirect	0.224 (0.078)	0.060 (0.033)	0.166 (0.070)	0.226 (0.078)
(iii) Total	0.364 (0.092)	0.096 (0.040)	0.270 (0.088)	0.366 (0.091)
(iv) Ratio (Indirect/Direct)	1.612 (0.763)	1.642 (1.069)	1.602 (0.857)	1.612 (0.761)
B: Scheduled News I	Version I	Version II		
	Combined	US	German	Combined
(v) Total	0.063 (0.036)	0.067 (0.037)	0.027 (0.023)	0.093 (0.046)
(vi) Ratio (Indirect/Direct)	1.399 (1.060)	1.223 (0.893)	1.589 (1.945)	1.317 (0.855)
C: Scheduled News II	Version I	Version II		
	Combined	US	German	Combined
(vii) Total	0.038 (0.034)	0.018 (0.017)	0.024 (0.028)	0.042 (0.037)
(viii) Ratio (Indirect/Direct)	1.249 (1.040)	1.646 (2.486)	1.105 (1.288)	1.303 (1.093)

Notes: The table reports elements of the variance decomposition for price changes implied by the GMM estimates of the daily models. Rows (i) - (iv) report estimates of $E[\Sigma_e^2(A_t)]/Var(\Delta p_t)$, $\alpha^2 E[\Sigma_u^2(A_t)]/Var(\Delta p_t)$, $(E[\Sigma_e^2(A_t)] + \alpha^2 E[\Sigma_u^2(A_t)])/Var(\Delta p_t)$ and $\alpha^2 E[\Sigma_u^2(A_t)]/E[\Sigma_e^2(A_t)]$. Under the Version I heading, the estimates use $A_t = A_t^{ALL}$. Under the US, German and Combined headings of Version II, A_t equals A_t^{US} , A_t^G and A_t^{ALL} . Panels B and C respectively report estimates using the number of scheduled news items, and the absolute forecast error for scheduled news. Standard errors computed from the estimated asymptotic distribution of the GMM estimates are reported in parentheses.

model estimates: $\alpha^2 E[\Sigma_u^2(A_t^{US}, A_t^G)]/E[\Sigma_e^2(A_t^{US}, A_t^G)]$. As the table shows, the contribution of news via the indirect channel is roughly 60 percent larger than the contribution via the direct channel. These estimates clearly indicate that the indirect effects of news operating via order flow are an important component of price dynamics.

Panels B and C of Table 7 are valuable for evaluating the robustness of our finding that indirect effects

are relatively important. The panels verify that the same conclusion arises across quite different definitions of news. In panel B we report elements of the unconditional price variance derived from Versions I and II of our daily model estimated with scheduled news only. This set includes US scheduled news concerning Non Farm Payroll, Durable Goods, Trade Balance, and Unemployment Claims, and all the German news items covered by MMS (see Section 2 above). As the table shows, estimates from Versions I and II of this modified model imply that scheduled news contributes between 6 and 10 per cent to the unconditional variance of prices. While these estimates are only about one quarter the size of their counterparts based on the full spectrum of news in panel A, they are still somewhat larger than the contribution implied by event studies. The ratio statistics reported in row (vi) suggest why this might be so. Here we see that, in sharp contrast to textbook models, scheduled news affects price more through the indirect channel than directly. It appears that scheduled news triggers trading that reveals incremental dispersed information.

As a further robustness check, we re-estimated the model using same set of scheduled news in a different way. The form of the model remains unchanged except that A_t^{US} and A_t^{G} in (21) now denote the sum of absolute standardized forecast errors for scheduled US and German news on day t . The forecast errors are computed from the surveys conducted by MMS, and the errors for each series are standardized using the standard deviation of the errors from January 1993 to December 1999.²¹ Panel C of Table 7 reports elements of the price variance derived from estimates of the daily model using this measure of scheduled news arrival. Here we see that scheduled news contributes about 4 percent to the variance of prices, while the ratios of indirect to direct effects remain above one.

To summarize, the results in Table 7 show that both scheduled and non-scheduled news contribute to the variance of the price changes in our sample. Our results also indicate that news items generally contain both common-knowledge information that is directly reflected in prices, and dispersed information that indirectly affects prices via its impact on order flow. Furthermore, all our model estimates imply that this indirect channel is at least as important as the direct channel in linking news to price changes.

5 Conclusion

This paper extends past work on FX prices and public news in three main ways. We address the presence of an indirect channel through which public news affects prices. Second, we use heteroskedasticity in order flow and price for identification, à la Rigobon and Sack (2004), rather than the more common event-study approach. Third, our methodology exploits the full set of macro news events piped into FX trading desks.

Our analysis of intraday data shows that order flow contributes more to changing FX prices in the period immediately following the arrival of news than at other times. This evidence pointing to the importance of the indirect channel is supported by our daily analysis: roughly two-thirds of the effect of macro news on FX

²¹Love and Payne (2004) construct a similar aggregate measure except that they “sign” each forecast error according to the direction of its theoretically predicted exchange rate effect. The latter adjustment is unnecessary here because our aim is to identify changes in the flow of macro news rather than to identify the directional influence of scheduled news on FX prices.

prices is transmitted via order flow, the remainder being the direct effect of news. With both the direct and indirect channels operating, we estimate that macro news accounts for 36 percent of total FX price variance in daily data. Given that daily prices are very nearly a martingale, this finding implies that macro news is far larger contributor to longer term price variation than previously thought.

Our daily results speak directly to the question, What drives order flow? The analysis in Evans and Lyons (2002a) splits total daily DM/\$ price variation into two parts: about 60 percent is due to order flow and about 40 percent is due to other factors. The results in Table 7 shed light on both of these parts. They suggest that order flow's 60 percent breaks roughly into one-third (20 percent) that is induced by macro news and two-thirds (40 percent) that is not news induced. Put differently, macro news accounts for about one-third of the variance of interdealer order flow in our sample. The 40 percent of total price variation due to other factors breaks into about one-third 15 percent from the direct effect of macro news and two-thirds (25 percent) that remains unaccounted for.

Finally, let us offer a wider perspective on our results. Inherent in current macro models is the view that price-setting marketmakers observe macro news, calculate the price implication, and instantly adjust all their FX prices by the same amount. Our results suggest that this is over-simplified. Rather, they suggest a model in which marketmakers observe macro news but have little idea how to interpret it, or how the rest of the market will interpret it. Instead, they wait to observe the trades induced and set their prices and expectations based on the interpretations embedded therein. (This view is consistent with the findings of Evans and Lyons 2005 that FX order flow conveys incremental information useful for forecasting macro variables.) Models with this richer informational structure may offer new insights into many of the long-standing puzzles concerning the behavior of FX prices.

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: Evidence from the U.S. Treasury Market, *Journal of Financial and Quantitative Analysis*, 36: 523-543.
- Berry, T., and K. Howe (1994), Public information arrival, *Journal of Finance*, 49: 1331-1346.
- Bollerslev, T., J. Cai, and F. Song (2000), Intraday periodicity, long-memory volatility, and macroeconomic announcement effects in the U.S. treasury bond market, *Journal of Empirical Finance*, 7: 37-55.
- Brandt, M., and K. Kavajecz (2004), Price discovery in the U.S. treasury market: The impact of order flow and liquidity on the yield curve, *Journal of Finance*, 59: 2623-2654.
- Cao, H., M. Evans, and R. Lyons (2006), Inventory information, *Journal of Business*, 79: 325-364.
- Carlson, J., and M. Lo (2004), One minute in the life of the DM/\$: Public news in an electronic market, *Journal of International Money and Finance*, forthcoming.
- 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. (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.
- Evans, M., and R. Lyons (2005), Exchange rate fundamentals and order flow, typescript, U.C. Berkeley, September.
- Fleming, M., (2003), Measuring treasury market liquidity, *Federal Reserve Bank of New York Economic Policy Review*, 9: 83-108.
- Fleming, M., and E. Remolona (1997), What moves the bond market? *Federal Reserve Bank of New York Economic Policy Review*, 3: 31-50.
- Fleming, M., and E. Remolona (1999), Price formation and liquidity in the U.S. treasury market, *Journal of Finance*, 54: 1901-1915.
- 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 (2005), Currency returns, intrinsic value, and institutional-investor flows, *Journal of Finance*, 60: 1535-1566.
- 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. (2004), Economic news and the impact of trading on bond prices, *Journal of Finance*, 59: 1201-1234.
- Hakkio C., and D. Pearce (1985), The reaction of exchange rates to economic news, *Economic Inquiry*, 23: 621-635.
- Hansen, L. (1982), Large sample properties of generalized method of moments estimators, *Econometrica*, 50: 1029-1054.
- Hardouvelis, G. (1988), Economic news, exchange rates, and interest rates, *Journal of International Money and Finance*, 7: 23-25.
- Hasbrouck, J. (1991), Measuring the information content of stock trades, *Journal of Finance*, 46: 179-207.
- Hayashi, F. (2000) *Econometrics*, Princeton University Press, Princeton NJ.
- Huang, R., J. Cai, and X. Wang (2002), Information-based trading in the interdealer market, *Journal of Financial Intermediation*, 11: 269-296.

- 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.
- Kandel, E., and N. Pearson (1995), Differential interpretation of public signals and trade in speculative markets, *Journal of Political Economy*, 103: 831-872.
- Kim, O., and R. Verrecchia (1994), Market liquidity and volume around earnings announcements, *Journal of Accounting and Economics*, 17: 41-67.
- Kim, O., and R. Verrecchia (1997), Pre-announcement and event-period information, *Journal of Accounting and Economics*, 24: 395-419.
- 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.
- Love, R., and R. Payne (2004), Macroeconomic news, order flows, and exchange rates, typescript, London School of Economics.
- Lyons, R. (2001), *The Microstructure Approach to Exchange Rates*, MIT Press.
- 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.
- 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.
- Payne, R. (2003), Informed trade in spot foreign exchange markets: An empirical analysis, *Journal of International Economics*, 61: 307-329.
- Potter, S. (1999), Non-linear time series modelling: An introduction, typescript, Federal Reserve Bank of New York.
- Rigobon, R., and B. Sack, (2004), The impact of monetary policy on asset prices, *Journal of Monetary Economics*, 51: 1553-1575.
- Rime, D. (2000), Private or public information in foreign exchange markets? An empirical analysis, typescript, University of Oslo, March.