

Myths about the Micro Approach to Exchange Rates

Martin D. D. Evans*

Richard K. Lyons

Draft: 18 June 2006

Abstract

We address four common misconceptions about micro-based research on exchange rates: (1) public news arrivals account for most exchange rate variation; (2) allocative trades do not convey information; (3) order flow is easy to measure; and (4) transactions obviously drive prices. Though few people subscribe to all four, most people subscribe to at least one. Dispelling these misconceptions should stimulate progress in this area.

Correspondence

Richard K. Lyons
Haas School of Business, UC Berkeley
Berkeley, CA 94720-1900
Tel: 510-642-1059, Fax: 510-643-1420
lyons@haas.berkeley.edu
faculty.haas.berkeley.edu/lyons

* Respective affiliations are Georgetown University and NBER, and UC Berkeley and NBER. We thank Mark Seasholes for comments and the National Science Foundation for financial assistance, which includes funding for a clearinghouse for recent micro-based research on exchange rates (at faculty.haas.berkeley.edu/lyons and at georgetown.edu/faculty/evansm1).

Myths about the Micro Approach to Exchange Rates

This paper addresses the recent micro-based literature on exchange rates, focusing specifically on four common misconceptions. Most people subscribe to at least one of them. The misconceptions are the following: (1) public news arrivals account for most exchange rate variation; (2) allocative trades do not convey information; (3) order flow is easy to measure; and (4) transactions obviously drive prices. The broader issues raised by these misconceptions are fundamental, namely, how public information gets into prices, how private information gets into prices, identifying information in trades, and synthesizing micro and macro views on price formation.

The following section presents the facts against each of these misconceptions. Section 2 addresses the directions that micro-based empirical work has gone in response to the facts. Section 3 concludes.

1. Four Misconceptions

Misconception 1: Public news accounts for most exchange rate variation.

What fraction of exchange rate movement is due to arrival of public macroeconomic news? 30 percent? 50 percent? 70 percent? Public news arrivals account for a disturbingly low fraction of total exchange rate variation – less than 5 percent. Despite scores of papers written on the link between exchange rates and news, this core fact remains.

This extreme result is not fully appreciated for a simple reason: the relevant papers are event studies, so their focus is on explanatory power within event windows, rather than explanatory power across full samples. Consider some important papers. A good example is Klein (1991), which presents analysis at the daily frequency. He regresses forex price changes on trade-balance news and finds that news explains about 40 percent of price changes on those days. 40 percent is quite significant. But remember that trade balance news arrives once per month, so roughly 95 percent of forex 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 the total price variation accounted for is less than 3 percent. Another fine paper in this literature is the recent paper by Andersen et al. (2003). They also find impressive R^2 statistics within their event windows, in this case, intraday windows. But as they note (p. 50), summing the amount of time across 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 conservative assumptions, their findings imply that news accounts for no more than 2 percent of the total price variation. We do not mean to suggest that these authors, or any others in the literature, have been disingenuous in presenting their findings. Rather, many readers have simply failed to put the R^2 statistics in broader perspective.

Though return volatility is not constant over time, this does not alter the message that the explanatory power of news is essentially zero. Our Klein (1991) example could include two adjustments in this respect: currency price volatility over weekends is not zero, which lowers his overall explanatory power (since his event windows never include weekends); but announcement days tend to have higher volatility than non-announcement days, which raises his explanatory power. Even with these adjustments, overall explanatory power remains below 5 percent. For Andersen et al. (2003), earlier work by these authors shows that the Employment Report has the largest impact on instantaneous variance, increasing it by a factor of ten (Andersen and Bollerslev 1998). To take an extreme case: If all announcements had the same factor-of-ten effect, and their within-event-window R^2 statistics were all one, news would still only account for 2 percent of total exchange rate variation. In fact, the R^2 statistics in Andersen et al. (2003) are generally below 50 percent (their Table 2), so an overall explanatory power of 2 percent is indeed an upper bound.

Misconception 2: Allocative trades do not convey information.

Information models of trading establish a conceptual dichotomy between trades motivated by speculation and trades motivated by allocation. Trades motivated by speculation convey information (Kyle 1985, Glosten and Milgrom 1985, Grossman and Stiglitz 1980). Trades motivated by allocation—whether due to endowment risk or liquidity needs—are noise. Informed agents know that they have information and

trade with the intent to profit from their information. Allocatively motivated trades can affect price, either because they are mistaken for informative trades or because they affect risk premia; but they do not convey fundamental payoff information.

Though this dichotomy has been useful for sharpening thinking, it has sharpened it a bit too much, in the sense of being taken too literally. As a theoretical matter, we offer a simple example that illustrates how allocatively motivated trades convey information. Consider a small, non-financial corporation that solves its optimization problem and realizes that it needs to trade foreign exchange, say because export proceeds are unexpectedly high this period, or because a burst in earnings in an overseas subsidiary needs to be hedged. The firm is trading forex for allocative reasons. It does not enter the forex market believing that it is informed. Nor is it informed, in isolation, in any measurable sense. But when this firm's trade is aggregated with those of other firms, the aggregate conveys information about the state of the current macroeconomy that might otherwise be reflected in official statistics months later, or not at all. Because these allocatively motivated trades involve the economy as a whole, it is in fact quite natural that they should convey deep economic information when aggregated. Yet, a person who takes a Kyle, Glosten-Milgrom, or Grossman-Stiglitz model too literally could easily dismiss the forex market as one to which information models do not apply.

Beyond this theoretical point, there is now empirical evidence that the trades of non-financial corporations—which most people view as allocatively motivated—convey fundamental information. For example, these trades convey information useful for forecasting exchange rates (Evans and Lyons 2005a) and for forecasting fundamental variables like real GDP and inflation (Evans and Lyons 2005b). Other markets such as government bond markets also share the forex market's focus on economy-wide information; it will not be surprising to find that allocative trades in these other markets also convey fundamental information.

Misconception 3: Order flow is easy to measure.

Order flow is central, both theoretically and empirically, for solving the identification problem. The relevant problem in this context is identifying the information in trades. Accordingly, measuring order flow correctly is of enormous significance. Though every trade involves a buyer and seller, generally the two sides are not the

same: one side has experienced a change in circumstances, perhaps new information, that causes a shift in its demand curve; whereas the other side is induced by price movement to take the other side of the trade, either as a dealer moving along her supply curve, or as an end-user who sees enough return to provide liquidity. The shifting demand curve embeds the information. Price-induced movement along a demand curve is not new information (so long as the slope of demand curve is known, which is very generally assumed). Order flow, as signed transaction flow, is our means of identifying the informative side, both conceptually and econometrically.

There are many difficulties in measuring order flow, including those arising from limit-order trading and those arising from the need to sign order flow from prices (Lee and Ready 1991). As an empirical matter, perhaps the purest—i.e., least prone to error—measures of order flow in the forex market come from direct interdealer trading. These are trades that occur bilaterally, with one dealer notifying another electronically that she would like a quote. The quote is provided electronically, and is typically available for a two-way decision for a few seconds. In data such as this, the aggressor is clear, i.e., the side whose demand curve has shifted is clear, as is the liquidity-providing side. Things are less clear in the context of limit-order trading. Though theory in this area assumes that limit orders come from the liquidity providers, so that order flow can be signed from the liquidity-taking market orders, this is a stretch empirically. There are many participants in the market who have experienced a shift in circumstances—a shift in their demand curve—who seek to reduce their transaction costs by trading into the position via limit orders. This results in trades being coded by the direction of the market order when in reality both sides reflect shifting demand curves. Measurement error from signing flows using prices a la Lee and Ready (1991) is well covered elsewhere in the literature.

The main point we want to make here goes beyond these well understood challenges in measuring order flow, and addresses instead a conceptual shortcoming that has appeared in the literature. This shortcoming is significant enough that it can lead to flawed analytical design—approaches that lose their ability to discriminate between hypotheses while on the surface appearing to have the power to do so. For example, this shortcoming can lead a researcher to mistakenly conclude that

there is no information in trades when in fact there is significant information in trades, but identification has not been achieved.

A simple example clarifies what we have in mind. It is a fact that forex dealers do not provide overnight liquidity—they generally do not take overnight positions. An important implication is that end-users, the other segment of trader in this market, are clearing this market among themselves at daily and lower frequencies. Some of these end-users are “push” traders in the sense of shifting demand curves above. Some are “pull” traders in the sense that they are pulled into the market by price movements, thereby playing a liquidity-providing role.¹ The key, then, is that one cannot identify news econometrically by measuring end-user transaction directions at the aggregate level. Indeed, to a first approximation, aggregate end-user “order flow” needs to sum to zero for market clearing. In this setting then, which is one in which the push trades are indeed moving prices, representative samplings from aggregate end-user order flow will have no explanatory power because it contains push and pull trades in equal measure (Froot and Ramadorai 2004).

There are (at least) two solutions to this identification problem. One is to build structural models to identify the information in flows. One such structural model is provided in Evans and Lyons (2002a). That model shows that the difficulty of identifying information in end-user flows can be solved by focusing on interdealer flows. In that model, the informative push trades are manifested in interdealer trading differently than the liquidity-providing pull trades, allowing the econometrician to disentangle them and test for information in trades in a robust way. A second solution is to partition the full aggregate of end-user flows into categories that separate, or at least begin to separate, the push trades from the pull trades. Of course, this too is a challenging task. But there has been some success in finding disaggregations that capture enough heterogeneity to provide statistical power (Rime et al.)

Misconception 4: Transactions obviously drive prices.

Market practitioners believe that transactions drive prices, and drive them persis-

¹ This nomenclature was introduced by Sager and Taylor (2005).

tently. This is as true in currency markets as it is in equity markets. On the academic side, researchers in microstructure finance have documented this in many ways and across many markets (Madhavan 2000, Lyons 2001). But researchers in economics more generally often do not share this view. Failing to recognize this disconnect risks leaving the micro-based literature on exchange rates “shouting in a jug.”

An example should help. Ken Rogoff, former Chief Economist at the IMF, recently wrote that,

“I still do not find compelling the notion that market absorption of asymmetric information (in the hands of individual traders) can have more than a very transitory impact on the market for major currency cross rates.” (Rogoff 2002)

This is a natural view, in accord with every textbook treatment of exchange rates. Textbook exchange rate models do not include any types of information that need to be aggregated by the market. Instead, information relevant to exchange rates is public and revealed publicly to the whole market simultaneously, either through official statistics (e.g., inflation) or via observable market variables (e.g., interest rates). It is well known that as an empirical matter, however, that the power of these macro models for explaining monthly or quarterly exchange rate movements among major currencies is essentially zero (Meese 1990, Frankel and Rose 1995). The view that trades do not convey information appears to be based on theoretical premises rather than the presence of a compelling empirical alternative.

In contrast, the empirical case that order flow has persistent impact on exchange rates is quite strong. For example, when interdealer order flow is measured precisely, it explains 40-80% of daily movements in major currency cross rates (Evans and Lyons 2002a,b). The word “daily” is important here: at the daily frequency, exchange rates are very nearly a martingale, as one would expect in a low friction speculative market. It is not possible in an efficient market for order flow to explain so much movement and have impact that is only transitory at the same time. Other evidence of persistent impact takes many forms, including impulse responses in systems with returns and order flow (Payne 2003), cointegration between exchange rate prices and cumulative order flow (Bjonnes and Rime 2003,

Killeen et al. 2006), and evidence that flows convey information about macro fundamentals (Evans and Lyons 2005b).

Despite the empirical record in favor of a persistent role for order flow in determining exchange rates, and the empirical record against more traditional macro approaches—both as monthly regression models and as event studies—the power of exchange rate theory remains strong in conditioning macroeconomic views. It is possible that some synthesis may emerge with recent findings that flows serve as a means by which information about the state of the macroeconomy is impounded in the exchange rate. Alternatively, synthesis may emerge in the direction of consensus that exchange rate movements have more to do with stochastic discount factors than they do with expected payoffs, as appears to be the case in asset pricing more generally. In the end, there is no getting around that there is a larger conceptual barrier to viewing markets this way than micro-oriented researchers realize.

2. Empirical Responses to the Facts

For misconception 1—the view that public news accounts for most exchange rate variation—we have already seen that the direct effects of news arrivals on exchange rates are disappointingly negligible. It is natural to ask whether news might affect exchange rates indirectly, i.e., via transactions that are induced by news. In fact, recent studies show that macroeconomic news arrivals stimulate transactions in the foreign exchange market, and that most of these induced transactions occur outside of traditional event windows (e.g., Love and Payne 2005, Evans and Lyons 2005a,b). Once the effects of these news-induced transactions are accounted for, the contribution of macro news in total price variation is an order of magnitude higher, in the 33–50 percent range. Our preferred interpretation of these results is that price setters in these markets, i.e., the dealers, do not find the implications of macro news for exchange rates easy to evaluate. Instead, they look to end-users in the market for the processing of information in announcements. The induced trades of the end-users, which are the signals of their processing conclusions, are typically not immediate. The picture appears to be much more nuanced than is suggested by the common-knowledge, rational expectations adjustment to news that is as a conceptual matter to us all.

For misconception 2—that allocative trades do not convey information—we have already referenced empirical evidence that allocatively motivated trades of non-financial corporations convey information about fundamentals like real GDP and inflation. Research is ongoing on whether allocative trades in other markets such as government bond markets or aggregate equity markets also share this property. They certainly share the forex market feature that what is relevant is information that is economy-wide. Still, the view implicit in this empirical strategy is that the information conveyed by allocative trades is limited to payoff information. It may be the case that allocatively motivated trades are more deeply related to market-wide discount factors than to expected future cash flows. Indeed, it is hard to imagine allocative flows at the aggregate level not having effects on discount factors. There is much room yet for further work on this front (see, e.g., Hasbrouck and Seppi 2001, Cao et al. 2006).

For misconception 3—that order flow is easy to measure—the state of affairs in foreign exchange trading are getting more challenging. The reason is that most interdealer trading now takes place through interdealer limit-order books rather than through direct bilateral trading. Recall that the latter allows quite precise discrimination between the liquidity taking side of a trade and the liquidity providing side. Until about ten years ago, the bulk of interdealer trading was done directly. Now roughly 90 percent occurs through limit-order books, principally systems provided by EBS and Reuters. The good news is that consolidated order flow data sets are likely to be increasingly available due to the consolidation of the interdealer market. The bad news is that without more guidance from limit-order theory, empiricists are inevitably working with less powerful measures of the information in trades. For consumers of this literature, this is an important development: it will be hard to disentangle whether weak results are actually describing markets or instead due to weak data.

For misconception 4—that transactions obviously drive prices—empiricists may have to take a different tack. Again, the issue is less one of piling on more empirical results than of establishing conceptual frameworks that more effectively bridge the gap between modern macroeconomics and micro approaches to security

price determination. One example is the model in Evans and Lyons (2005b). That framework is the workhorse framework in international macroeconomics, namely, dynamic general equilibrium (DGE) modeling. Recent developments in DGE modeling include greater articulation in terms of preference structures (e.g., habit persistence) and in terms of production structures (e.g., the supply side of open-economy macro models a la Obstfeld and Rogoff 1995). The bridge is likely to come from greater articulation of these models in a different direction: information structures. DGE modeling generally assumes common knowledge information. The model in Evans and Lyons (2005b) relaxes this and shows how information is aggregated via trading in speculative markets, in their case, the foreign exchange market. The complexity of the model is likely to appear unwarranted to many researchers from the micro-based tradition. If micro-tradition models were more successful in speaking to macroeconomists, we might agree. At this time, it is hard to envision greater micro-macro synthesis without greater commonality in frameworks.

3. Conclusions

The four misconceptions addressed here are prevalent. Most people subscribe to at least one. They are certainly not the only misconceptions associated with micro-based exchange rate research. We chose to focus on these because the broader issues they raise are fundamental. The first misconception raises the question of how public information gets into price. Event studies of how this works find remarkably little evidence that news arrivals do the work that textbook models imply that they should. This result is a significant challenge to the field. Putting an end to this misconception should help to move ahead more constructively.

The second misconception raises the question of how private information gets into price. In traditional exchange rate economics, this is a non-question: there is no private information in forex markets. If private information is understood to mean concentrated inside information, then we would agree, this is not the right way to think about the major markets. But the category of information types other than public information includes much, much more than concentrated inside information.

It includes any type of information that needs to be aggregated by markets, sometimes referred to as “dispersed” information. Dispersed information may be about fundamental payoffs. But it also may be about economy-wide discount factors. The point is that “private” information in the forex market is indeed relevant in this sense. And this type of information need not be associated with agents who participate in this market with the knowledge that they have superior information. Rather, it can be communicated by trades that are allocative—not motivated by a conscious effort to exploit information at the micro level.

The third misconception raises the question of how to identify the information in trades. This is made easy in theory, typically by assuming that the trades of liquidity providers—for example in the form of limit-order traders—do not convey any information. As an empirical matter, matters are much more complex. There is a lot at stake: it is foreseeable that researchers will mistakenly conclude that there is little information in trades, when in fact identification has not been achieved. Moving ahead constructively on this issue is likely to require more theoretical participation on the part of empiricists.

The fourth misconception raises the question of intellectual communication and the importance of common frameworks for synthesizing views on price formation. Some would argue that this is the least significant of the misconceptions. Perhaps, at least when viewed in the narrow as addressing whether micro-based research on exchange rates will influence macroeconomics. But there is a larger issue raised here: what is the role of information imperfections in modern macroeconomics? Can the puzzles in macro be addresses satisfactorily via greater articulation of preference and production structures? Or might the information dimension of these models need great articulation? A richer micro-macro dialogue might very well aspire to this larger goal.

References

- Andersen, T., and T. Bollerslev, 1998, Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies, *Journal of Finance*, 53: 219-266.
- 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.
- Bjonnes, G., and D. Rime (2005), Dealer behavior and trading systems in foreign exchange markets, *Journal of Financial Economics*, 75: 571-605.
- Bjonnes, G., D. Rime, and H. Solheim (2005), Liquidity provision in the overnight foreign exchange market, *Journal of International Money and Finance*, 24: .
- 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.
- 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), Informational integration and FX trading, *Journal of International Money and Finance*, 21: 807-831.
- Evans, M., and R. Lyons (2005a), Meese-Rogoff redux: Micro-based exchange rate forecasting, *American Economic Review*, June 2005.
- Evans, M., and R. Lyons (2005b), Exchange rate fundamentals and order flow, typescript, U.C. Berkeley, May.
- Fleming, M., and E. Remolona (1999), Price formation and liquidity in the U.S. treasury market, *Journal of Finance*, 54: 1901-1915.
- Frankel, Jeffrey, and Andrew Rose, 1995, Empirical research on nominal exchange rates, in G. Grossman and K. Rogoff (eds.), *Handbook of International Economics*, Elsevier Science: Amsterdam, 1689-1729.
- Froot, K., and T. Ramadorai (2005), Currency returns, intrinsic value, and institutional-investor flows, *Journal of Finance*, 60: 1535
- 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.
- Grossman, S., and J. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review*, 70: 393-408.
- Hasbrouck, J. (1991), Measuring the information content of stock trades, *Journal of Finance*, 46: 179-207.
- Hasbrouck, J., and D. Seppi (2001), Common factors in prices, order flows, and liquidity, *Journal of Financial Economics*, 59: 383-411.
- Killeen, W., R. Lyons, and M. Moore (2006), Fixed versus flexible: Lessons from EMS Order Flow, *Journal of International Money and Finance*.
- 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.
- Lee, C., and M. Ready, 1991, Inferring trade direction from intradaily data, *Journal of Finance*, 46, 733-746.
- Love, R., and R. Payne (2005), 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.
- Meese, R., 1990, Currency fluctuations in the post-Bretton Woods era, *Journal of Economic Perspectives*, 4: 117-134.
- Meese, R., and K. Rogoff, 1983a, Empirical exchange rate models of the seventies, *Journal of International Economics*, 14: 3-24.
- Obstfeld, M. and K. Rogoff (1995), Exchange Rate Dynamics Redux, *Journal of Political Economy*, 103: 624-660.
- Payne, R. (2003), Informed trade in spot foreign exchange markets: An empirical analysis, *Journal of International Economics*, 61: 307-329.
- Rogoff, K. (2002), in *Economic Policy*, discussion of papers on transaction costs and the euro as an international currency.
- Sager, M., and M. Taylor (2005), Under the microscope: The structure of the foreign exchange market, Warwick University typescript.