Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market

Alain Chaboud Benjamin Chiquoine Erik Hjalmarsson Clara Vega*

June 13, 2011

Abstract

We study the impact of algorithmic trading in the foreign exchange market using a high-frequency dataset representing a majority of global interdealer trading in three major currency pairs, euro-dollar, dollar-yen, and euro-yen, from 2003 through 2007. We find that human-initiated trades account for a larger share of the variance in exchange rate returns than computer-initiated trades: humans are still the "informed" traders. There is some evidence, however, that algorithmic trading contributes to a more efficient price discovery process via the elimination of triangular arbitrage opportunities and the faster incorporation of macroeconomic news surprises into the price. We also show that algorithmic trades tend to be correlated, indicating that computer-driven strategies are not as diverse as those used by human traders. Despite this correlation, we find no evidence that algorithmic trading causes excess volatility. Furthermore, the amount of algorithmic activity in the market has a small, but positive, impact on market liquidity.

JEL Classification: F3, G12, G14, G15.

Keywords: Algorithmic trading; Liquidity provision; Price discovery; Private information.

---

*Chaboud and Vega are with the Division of International Finance, Federal Reserve Board, Mail Stop 43, Washington, DC 20551, USA; Chiquoine is with the Investment Fund for Foundations, 97 Mount Auburn Street, Cambridge MA 02138, USA; Hjalmarsson is with Queen Mary, University of London, School of Economics and Finance, Mile End Road, London E1 4NS, UK. Please address comments to the authors via e-mail at alain.p.chaboud@frb.gov, bchiquoine@tiff.org, e.hjalmarsson@qmul.ac.uk, and clara.vega@frb.gov. We are grateful to EBS/ICAP for providing the data, and to Nicholas Klagge and James S. Hebden for their excellent research assistance. We would like to thank Cam Harvey, an anonymous Associate Editor and an anonymous referee for their valuable comments. We also benefited from the comments of Gordon Bodnar, Charles Jones, Terrence Hendershott, Luis Marques, Albert Menkveld, Dagfinn Rime, Alec Schmidt, John Schoen, Noah Stoffman, and of participants in the University of Washington Finance Seminar, SEC Finance Seminar Series, Spring 2009 Market Microstructure NBER conference, San Francisco AEA 2009 meetings, the SAIS International Economics Seminar, the SITE 2009 conference at Stanford, the Barcelona EEA 2009 meetings, the Essex Business School Finance Seminar, the EDHEC Business School Finance Seminar, and the Imperial College Business School Finance Seminar. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

1 Introduction

The use of algorithmic trading, where computers monitor markets and manage the trading process at high frequency, has become common in major financial markets in recent years, beginning in the U.S. equity market in the late 1990s. Since the introduction of algorithmic trading, there has been widespread interest in understanding the potential impact it may have on market dynamics. While some have highlighted the potential for improved liquidity and more efficient price discovery, others have expressed concern that it may lead to reduced liquidity and excessive volatility, particularly in times of market stress. Despite the widespread interest, formal empirical research on algorithmic trading has been rare, primarily because of a lack of data in which algorithmic trades are clearly identified.[1]

This paper investigates whether algorithmic ("computer") trades and non-algorithmic ("human") trades have different effects on the foreign exchange market; it is the first formal empirical study on the subject in the foreign exchange market. We rely on a novel dataset consisting of several years (October 2003 to December 2007) of minute-by-minute trading data from Electronic Broking Services (EBS) in three currency pairs: the euro-dollar, dollar-yen, and euro-yen. The data represent a majority of spot interdealer transactions across the globe in these exchange rates. A crucial feature of the data is that, on a minute-by-minute frequency, the volume and direction of human and computer trades are explicitly identified, allowing us to measure their respective impacts at high frequency. Another useful feature of the data is that it spans the introduction and rapid growth of algorithmic trading in a market where it had not been previously allowed.

In a classical frictionless world, there is little reason why trading by computers and humans should have different impacts. In a more realistic setting, Foucault, Kadan, and Kandel (2009) model algorithmic traders as incurring lower monitoring costs than humans, which allows algorithmic traders to react more quickly to changes in the limit order book. In related work, Biais, Foucault, and Moinas (2010) analyze the effect that the interaction between fast and slow investors may have on market liquidity and the informational efficiency of prices. Both of these studies highlight the key feature of algorithmic trading, its advantage in speed over human trading. In practice, this advantage might be attributed to both the rapidity with which a computer can follow a particular decision rule and execute a given trading command, but also, as in both models, to the ability of computers to access and monitor vast amounts of information in a faster and more timely fashion than humans. In both of these models, algorithmic trading improves the speed at which certain types of information are incorporated into prices. However, the effect algorithmic trading

[1] A notable early exception prior to our study is the paper by Hendershott, Jones, and Menkveld (2011), who got around the data constraint by using the flow of electronic messages on the NYSE as a proxy for algorithmic trading. They concluded that algorithmic trading on the NYSE, contrary to the pessimists' concerns, likely caused an improvement in market liquidity. Subsequent to the first version of our study, Hendershott and Riordan (2009), Brogaard (2010), and Zhang (2010), have conducted empirical work on the subject.

has on liquidity is ambiguous. Foucault et al. (2009) argue that, when monitoring costs for market-takers decline, liquidity is consumed relatively more quickly than it is supplied and, as a consequence, the depth of the market decreases. In contrast, when monitoring costs for market-makers decline, the opposite effect occurs, resulting in an increase in depth. In Biais et al. (2010)'s model, the introduction of algorithmic trading has two opposing effects. The introduction of algorithmic traders increases the probability of finding a counterparty to trade with and therefore it can increase the depth of the market. But the introduction of faster and better informed algorithmic traders raises adverse selection costs for slower traders, leading some to leave the market. Depending on which effect dominates, the introduction of algorithmic trading may increase or decrease market liquidity.

A potential drawback of computer-based trading is that it must be fully rule-based, and therefore void of human judgements and instincts that cannot be pre-programmed. Although this limitation may have its upsides, such as avoiding emotional biases in trading decisions (e.g., Lo, Repin and Steenbarger, 2005), it may also limit the scope and diversity of possible algorithmic trading strategies. This is an often-cited concern in the press about algorithmic trading, where the fear is that the lack of diversity may lead to large blocks of similar trading instructions reaching the market at the same time, resulting in reduced liquidity and excess volatility. Khandani and Lo (2007, 2011), who analyze the large losses that occurred for many quantitative long-short equity strategies at the beginning of August 2007, highlight the possible adverse effects on the market of such commonality in behavior across market participants (algorithmic or not) and provide empirical support for this concern. From a theoretical perspective, Biais et al. (2010) argue that in times of stress, the lack of flexibility of algorithmic traders could cause them to submit aggressive orders that trigger excessive price movements which are later reversed. Also, if algorithmic traders behave *en masse* like the positive-feedback traders of DeLong, Shleifer, Summers, and Waldman (1990) or the chartists described in Froot, Scharfstein, and Stein (1992), it is conceivable that they might cause deviations of prices from fundamental values and thus induce excess volatility.

Our empirical analysis is motivated by the above considerations. We investigate whether algorithmic traders behave differently and have a different impact on the foreign exchange market than human traders. Given the novelty of our dataset and the lack of any prior empirical work on algorithmic trading in the foreign exchange market, we begin the study with summary statistics and document several stylized facts. First, computers tend to trade in smaller sizes than humans, which is consistent with computers incurring lower monitoring costs than humans and with the computers' ability to return to the market more quickly, with less risk of a movement in prices. Second, as in the studies on equity markets by Biais, Hillion, and Spatt (1995) and Hendershott and Riordan (2009), we find that trades of the same type (e.g., small buys by computers or large buys by humans) are more likely to follow each other than trades of two different types, which is consistent with strategic order splitting either for noninformational reasons (e.g., Bertsimas and Lo (1998)) or for informational reasons (e.g., Easley and O'Hara (1987)). Third, trades initiated by computers tend to be more serially correlated than those initiated by humans, which suggests that strategic order splitting is implemented relatively more via computer trading than via human trading. Fourth, the introduction and growth of algorithmic trading coincides with a dramatic reduction in triangular arbitrage opportunities and with a faster response of exchange rates to the release of macroeconomic data, in accordance with the models of Foucault et al. (2009) and Biais et al. (2010).

The more formal empirical analysis starts with the estimation of return-order flow dynamics in a structural vector autoregression (VAR) framework in the tradition of Hasbrouck (1991a). The VAR estimation provides an important insight: human order flow explains a substantially larger share of the variance in exchange rates than does computer order flow.[2] In other words, humans seem to be the better "informed" traders, in the sense that their order flow has a bigger impact on prices. This is in contrast to findings in equity market studies (e.g., Hendershott and Riordan, 2009), where computer order flow is typically found to have a larger impact on prices. The exception among our results is for the euro-yen currency pair, where we find that computers contribute about as much as humans to price discovery. The euro-yen is effectively the residual cross-rate in the three currency pairs we analyze, and is thus primarily determined by the triangular arbitrage relationship that must hold between the three pairs, rather than by direct price discovery. The VAR results suggest that, in this cross-rate, computers and humans are equally informed, which may reflect the fact that computers have a natural speed advantage in trading on triangular arbitrage opportunities.

We investigate next whether the trading strategies used by computers are more correlated than those used by humans. Since we do not have specific data on trading strategies, we indirectly infer the correlation among computer trading strategies from our trading data. The primary idea behind the correlation test that we design is that traders who follow similar trading strategies will trade less with each other than those who follow less correlated strategies. In a simple random matching model, the amount of trading between two types of traders is determined by the relative proportion of each type of trader in the market. Comparing the model's predictions to the realized values in the data, we find very strong evidence that computers do not trade with each other as much as the model would predict. There is thus evidence that the strategies embodied in the computer algorithms are indeed more correlated and less diverse than those used by human traders.

Finally, we analyze whether algorithmic trading causes market depth (liquidity) and/or "excess" volatility.[3] The amount of algorithmic trading as a fraction of total trading volume is likely endogenous with

[2] Order flow is defined as the net of buyer-initiated trading volume minus seller-initiated trading volume.

[3] Our definition of excess exchange rate volatility is volatility that cannot be explained by past volatility and contemporaneous macroeconomic news releases.

regards to variables such as liquidity and volatility, in the sense that there is a potential causal relationship going both ways. Standard ordinary least squares (OLS) estimation of the impact of algorithmic trading on market liquidity and volatility is therefore likely biased. In order to alleviate this problem, we rely on both Granger causality tests and instrumental variable (IV) estimation. In the IV estimation, we instrument for the level of algorithmic trading with a variable that measures the installed capacity for algorithmic trading: a monthly times-series of the share of trading floors equipped to trade algorithmically on the EBS system. To complement the IV results, we also estimate Granger causality tests using minute-by-minute data. Both estimation methods have their advantages and disadvantages. Taken together, we conclude that algorithmic trading has a slight positive effect on market liquidity in the euro-dollar and dollar-yen market. We also conclude that there is no systematic evidence to back the often-voiced concern that increased algorithmic trading leads to excessive volatility in any of the three currency pairs we analyze.

The paper proceeds as follows. Section 2 introduces the high-frequency data used in this study, including a short description of the structure of the market and an overview of the growth of algorithmic trading in the foreign exchange market over time. Section 3 presents some initial descriptive data analysis. Section 4 studies price discovery in a vector autoregressive framework. Section 5 analyzes whether computer strategies are more correlated than human strategies. In Section 6, we test whether there is evidence that the share of algorithmic trading in the market has a causal impact on market liquidity and volatility. Finally, Section 7 concludes. Some additional clarifying and technical material is found in the Appendix.

2 Data description

2.1 Market structure

Over our sample period, from 2003 to 2007, two electronic platforms processed the majority of global interdealer spot trading in the major currency pairs, one offered by Reuters, and one offered by Electronic Broking Services (EBS).[4] Both of these trading platforms are electronic limit order books. Importantly, trading in each major currency pair is highly concentrated on only one of the two systems. Of the most traded currency pairs (exchange rates), the top two, euro-dollar and dollar-yen, trade primarily on EBS, while the third, sterling-dollar, trades primarily on Reuters. As a result, price discovery for spot euro-dollar, for instance, occurs on the EBS system, and dealers across the globe base their customer and derivative quotes on that price. EBS controls the network and each of the terminals on which the trading is conducted. Traders can enter trading instructions manually, using an EBS keyboard, or, upon approval by EBS, via a computer di-

[4]EBS, which was previously owned by a group of foreign-exchange dealing banks, was purchased by the ICAP group in 2006. ICAP owns interdealer trading platforms and voice-broking services for most types of financial assets in a number of countries.

rectly interfacing with the system. The type of trader (human or computer) behind each trading instruction is recorded by EBS, allowing for our study.

The EBS system is an interdealer system accessible to foreign exchange dealing banks and, under the auspices of dealing banks (via prime brokerage arrangements), to hedge funds and commodity trading advisors (CTAs). As it is a "wholesale" trading system, the minimum trade size is 1 million of the "base" currency, and trade sizes are only allowed in multiple of millions of the base currency. We analyze data in the three most-traded currency pairs on EBS, euro-dollar, dollar-yen, and euro-yen.[5]

2.2 Price, volume, and order flow data

Our data consists of both quote data and transactions data. The quote data, at the one-second frequency, consist of the highest bid quote and the lowest ask quote on the EBS system in our three currency pairs. The quote data are available from 1997 through 2007. All the quotes are executable and therefore truly represent the market price at that instant.[6] From these data, we construct mid-quote series from which we can compute exchange rate returns at various frequencies. The transactions data, available from October 2003 through December 2007, are aggregated by EBS at the one-minute frequency. They provide detailed information on the volume and direction of trades that can be attributed to computers and humans in each currency pair. Specifically, each minute we observe trading volume and order flow for each of the four possible pairs of human and computer makers and takers: human-maker/human-taker (HH), computer-maker/human-taker (CH), human-maker/computer-taker (HC), and computer-maker/computer-taker (CC).[7]

Figure 1 shows, from 2003 through 2007, for each currency pair, the fraction of trading volume where at least one of the two counterparties is an algorithmic trader, i.e., $Vol(CH + HC + CC)$ as a fraction of total volume. From its beginning in late 2003, the fraction of trading volume involving algorithmic trading grows by the end of 2007 to near 60 percent for euro-dollar and dollar-yen trading, and to about 80 percent for euro-yen trading. As algorithmic trading represents only a small share of total volume in the early part of our sample, much of the upcoming analysis will focus on the trading data from 2006 and 2007. We will also consider a shorter and more recent sub-sample of data from September, October, and November of 2007,

[5] The euro-dollar currency pair is quoted as an exchange rate in dollars per euro, with the euro the "base" currency. Similarly, the euro is also the base currency for euro-yen, while the dollar is the base currency for the dollar-yen pair.

[6] In our analysis, we exclude data collected from Friday 17:00 through Sunday 17:00 New York time from our sample, as activity on the system during these "non-standard" hours is minimal and not encouraged by the foreign exchange community. Trading is continuous outside of the weekend, but the value date for trades, by convention, changes at 17:00 New York time, which therefore marks the end of each trading day. We also drop certain holidays and days of unusually light volume: December 24-December 26, December 31-January 2, Good Friday, Easter Monday, Memorial Day, Labor Day, Thanksgiving and the following day, and July 4 (or, if this is on a weekend, the day on which the U.S. Independence Day holiday is observed).

[7] The naming convention for "maker" and "taker" reflects the fact that the "maker" posts quotes before the "taker" chooses to trade at that price. Posting quotes is, of course, the traditional role of the market-"maker." We refer the reader to Appendix A1 for more details on how we calculate volume and order flow for these four possible pairs of human and computer makers and takers.

where the share of algorithmic activity is even greater.[8]

Figure 2 shows the evolution over time of the four different possible types of trades: $Vol(HH)$, $Vol(CH)$, $Vol(HC)$, and $Vol(CC)$, as fractions of the total volume. By the end of 2007, in the euro-dollar and dollar-yen markets, human to human trades, the solid lines, account for slightly less than half of the volume, and computer to computer trades, the dotted lines, for about ten to fifteen percent. In these two currency pairs, $Vol(CH)$ is often slightly higher than $Vol(HC)$, i.e., computers "take" prices posted by humans, the dashed lines, less often than humans take prices posted by market-making computers, the dotted-dashed lines. The story is different for the cross-rate, the euro-yen currency pair. By the end of 2007, there are more computer to computer trades than human to human trades. But the most common type of trade in euro-yen is computers trading on prices posted by humans. We believe this reflects computers taking advantage of short-lived triangular arbitrage opportunities, where prices set in the euro-dollar and dollar-yen markets, the primary sites of price discovery, are very briefly out of line with the euro-yen cross rate. Detecting and trading on triangular arbitrage opportunities is widely thought to have been one of the first strategies implemented by algorithmic traders in the foreign exchange market, which is consistent with the more rapid growth in algorithmic activity in the euro-yen market documented in Figure 1. We discuss the evolution of the frequency of triangular arbitrage opportunities in Section 3 below.

Summary statistics for the one-minute order flow data are shown in Table 1. This table contains two noteworthy features. First, order flow is positively serially correlated. We analyze this positive serial correlation in more detail in Section 3, where we estimate the conditional frequencies of various sequences of different types of trades. Second, the standard deviations of the various order flows differ by exchange rates, by type of taker and across maker/taker pairs; these differences will be important in the interpretation of the VAR analysis and variance decompositions that we present in Section 4.

3 Descriptive data analysis

3.1 Trade size

Equity market studies have found that computers often trade in smaller sizes (e.g., Hendershott, Jones and Menkveld (2011) and Hendershott and Riordan (2009)). This finding is consistent with the view that some computer algorithms are used to split large orders into smaller trades. It is also potentially consistent with computers incurring lower monitoring costs than humans, as the more frequent monitoring by computers may allow them to trade again sooner than humans, with less risk of a movement in price.

[8] December is omitted to ensure that potential year-end effects do not have an undue influence on the results, given the fairly short three-month sample period.

We show in Figure 3 average trade sizes across time for the period 2006 to 2007, broken down by who took the liquidity, human-taker ($HH+CH$), labeled as $H-Take$, or computer-taker ($HC+CC$), labeled as $C-Take$, in the first column of graphs, and by who provided the liquidity, human-maker ($HH+HC$), labeled as $H-Make$, and computer-maker ($CC+CH$), labeled as $C-Make$, in the second column of graphs. The third column shows the most disaggregated decomposition of trades: human-maker/human-taker (HH), computer-maker/human-taker (CH), human-maker/computer-taker (HC), and computer-maker/computer-taker (CC). As seen in the graphs, the average trade size for each category falls between 1 and 2 millions of the base currency. Regardless of whether computers are taking liquidity from the market or providing liquidity to the market, the average algorithmic trade size is smaller than the average human trade size. Interestingly, this fact is consistent with the findings in equity markets, even though the two market structures are very different and the variability in trade size in the interdealer foreign exchange market, with its large minimum trade size, is very small compared to the variability in trade size in the equity market. Another noteworthy feature of the graphs is that there is a clear ordering in trade size in the most disaggregated decomposition: on average HH trades are the largest, then CH trades, and then HC trades. CC trades, where computers are both providers and takers of liquidity, are on average the smallest.

3.2 Serial correlation of different types of trades

We next report some statistics on the first-order serial correlation of different types of trades, similar to those reported in Biais, Hillion, and Spatt (1995) and Hendershott and Riordan (2009). We classify trades by size (small or large), by the type of initiator (human or computer taker), and by whether the initiator is buying or selling the base currency. In each column of Table 2, we report the average conditional frequencies of various sequences of different types of trades, e.g., a small computer buy followed by a small human buy. Because computer participation increased greatly over our sample period, we report results for the three month sub-sample covering the months of September, October, and November of 2007, when algorithmic activity is largest in our sample. We note, though, that the findings described below are qualitatively similar when we use the full sample to estimate the conditional frequencies.

Each row in Table 2 sums up to 1 and can be interpreted as a probability vector. Eight different trade types are considered: "Human Small Buy" and "Human Small Sell" orders (denoted HSB and HSS, respectively), "Human Large Buy" and "Human Large Sell" orders (denoted HLB and HLS, respectively), and the corresponding computer trades (denoted CSB, CSS, CLB, and CLS, respectively). The exact definition of these trade categories based on our one-minute trading data are given in Appendix A2. The no trade event is denoted by "Z", and the row labeled "Unconditional" reports the average unconditional fraction of a given trade type. The row labeled "Likelihood ratio" allows us to compare the diagonal terms across columns. For example, the diagonal likelihood ratio of a small buy order from a human (HSB) following a small buy order from a human (equal to 1.28 in the EUR/USD exchange rate) is defined as

$$\frac{\Pr(HSB_t|HSB_{t-1})}{\Pr(HSB_t|HSB^C_{t-1})} = \Pr(HSB_t|HSB_{t-1}) \times \frac{\Pr(HSB^C_{t-1})}{\Pr(HSB_t|HSS_{t-1})\Pr(HSS_{t-1}) + ... + \Pr(HSB_t|Z_{t-1})\Pr(Z_{t-1})},$$

where HSB^C_{t-1} denotes the complement of HSB_{t-1}. The (diagonal) likelihood ratio for a given trade type thus captures the extent to which trades of a given type are likely to follow each other compared to the likelihood of observing such a trade type when a different type occurred in the previous period. Loosely speaking, the likelihood ratio can therefore be seen as a measure of the relative persistence of a given trade type and it allows for direct comparison of this persistence across different trade types.

Following Biais, Hillion, and Spatt (1995) and Hendershott and Riordan (2009), we highlight in bold the largest value in each column. Consistent with the findings in these two papers, the diagonal terms in Table 2 are in bold, which implies that trades of the same type are more likely to follow each other than trades of two different types. The succession of identical types of orders could reflect strategic order splitting, either for noninformational reasons (e.g., Bertsimas and Lo (1998)) or for informational reasons (e.g., Easley and O'Hara (1987)). Alternatively, He and Wang (1995) show that insiders tend to trade until their private information is incorporated into the market. It is also possible that traders of a given type react similarly, but successively, to the same events.

Table 2 shows that for the euro-dollar and dollar-yen currency pairs, the persistence of computer trades of a particular type is higher than the persistence of human trades of a particular type. For instance, in the euro-dollar currency pair, the likelihood ratio of a "Computer Small Buy" trade following another is 1.85, compared to a likelihood ratio of 1.28 for a "Human Small Buy" trade following another. This finding is consistent with the idea that strategic order splitting in the foreign exchange market is implemented more via computer trading than via human trading. In the cross-currency pair, the euro-yen, the opposite result holds, with human trades of a given type often more persistent than computer trades of the same type. This might again reflect the fact that computers are often used for triangular-arbitrage trading in the cross currency, with the trading patterns not likely to be as persistent given the transient nature of such opportunities. We next investigate whether the introduction of algorithmic traders is associated with a decrease in the number of triangular arbitrage opportunities.

3.3 The frequency of triangular arbitrage opportunities.

As discussed in the introduction, the models of Foucault et al. (2009) and Biais et al. (2010) predict that algorithmic trading improves the speed with which information is incorporated into prices. One obvious way for computers to have an impact on informational efficiency in the foreign exchange market is by detecting and taking advantage of triangular arbitrage opportunities, thus reducing the occurrence of such opportunities. Since our data contain second-by-second prices on three related exchange rates (euro-dollar, dollar-yen, and euro-yen), we can test the frequency with which these exchange rates are "out of alignment." More precisely, each second we evaluate whether a trader, starting with a dollar position, could profit from purchasing euros with dollars, purchasing yen with euros, and purchasing dollars with yen, all simultaneously at the relevant bid and ask prices. An arbitrage opportunity is recorded for any instance when such a strategy (and/or a "round trip" in the other direction) would yield a profit of two basis points or more.[9] The daily frequency of such opportunities is shown from 2002 through 2007 in Figure 4.

The frequency of arbitrage opportunities drops dramatically over our sample, with the drop being particularly noticeable around 2005, when the rate of growth in algorithmic trading is highest. On average in 2003 and 2004, the frequency of such arbitrage opportunities is about 0.1 percent (0.119 and 0.091 percent, respectively), one occurrence every 1,000 seconds. By 2007, at the end of our sample, the frequency has declined to 0.006 percent, less than one occurrence every 10,000 seconds. This simple analysis highlights the potentially important impact of algorithmic trading in this market. It is clear that other factors could have contributed to, or even driven, the drop in arbitrage opportunities, and the analysis certainly does not prove that algorithmic trading caused the decline. However, the findings line up well with the anecdotal (but widespread) evidence that one of the first strategies widely implemented by algorithmic traders in the foreign exchange market aimed to detect and profit from triangular arbitrage opportunities.

3.4 Price response to the release of macroeconomic data

Among the most recent developments in algorithmic trading, some algorithms now automatically read and interpret economic data releases, generating trading orders before economists have begun to read the first line. In this section, we document changes over time in the immediate response of the euro-dollar and dollar-yen exchange rates to the release of U.S. nonfarm payroll data, one of the most closely observed U.S.

[9] We conduct our test over the busiest period of the trading day in these exchange rates, from 3:00 am to 11:00 am New York Time, when all three exchange rates are very liquid. Our choice of a two-basis-point profit cutoff ($200 per $1 million traded) is arbitrary but, we believe, reasonable; the frequencies of arbitrage opportunities based on several other minimum profit levels show a similar pattern of decline over time. Note that even though we account for actual bid and ask prices in our calculations of profits, an algorithmic trader also encurs other costs (e.g., overhead, fees for the EBS service, and settlement fees). In addition, the fact that trades on the system can only be made in whole millions of the base currency creates additional uncertainty and implied costs. As an example, if a trader sells 2 million dollars for 1.5 million euros, the next leg of the triangular arbitrage trade on EBS can only be a sale of, say, 1 or 2 million euros for yen, not a sale of 1.5 million euros. Therefore, even after accounting for bid-ask spreads, setting a minimum profit of zero to detect triangular arbitrage opportunities is not realistic.

macroeconomic announcements.[10,11] We note, that the results described below are qualitatively similar when considering other important macroeconomic news releases.

We estimate the price response in the first minute following the data release, which always occurs precisely at 8:30 a.m. Eastern Time (ET). We estimate the responses to data releases in three separate sub-periods, 2002-2003, 2004-2005, and 2006-2007. In each sub-period, we estimate the second-by-second cumulative response of the exchange rate to the 24 data releases by regressing the exchange rate return, $100 \times \ln(p_{t+i}/p_{t-1})$, on the nonfarm payroll surprise (actual announcement minus market expectation); p_{t-1} is the price of the asset one second before the data release and p_{t+i} is the price of the asset i seconds after the release, $i = 1, ..., 60$. Figure 5 plots the estimated slope coefficients over the 60 seconds following the data release. To compare responses across exchange rates, we standardize the coefficients so that they are equal to one on average in the 60-second period after the data release (i.e., we divide each coefficient by the average magnitude of the coefficients). In each figure we also report, in a box, the average absolute deviation of the coefficients from one, as a measure of the amount of under- or over-reaction of the prices in the immediate period after the data release.

For both exchange rates, the price response to the nonfarm payroll announcement is clearly faster in 2006-2007 than in 2002-2003, before the introduction of algorithmic trading. Not only is the initial response faster—the line is steeper in the first few seconds after the data release—but, as shown by the average deviation of the coefficients from one, the price also stabilizes more firmly around its new "equilibrium" level in the latter sample period, after algorithmic trading has become prevalent. This finding fits well the models of Foucault et al. (2009) and Biais et al. (2010), which predict that algorithmic trading should increase the speed with which information is incorporated into prices, contributing to market efficiency. As in the case of triangular arbitrage opportunities, this evidence is, again, only suggestive of the impact of algorithmic trading, as we do not control for changes over time in other factors that might influence the speed with which information is incorporated in prices.

[10] Andersen and Bollerslev (1998), among others, refer to the nonfarm payroll report as the"king" of announcements because of the significant sensitivity of most asset prices to its release.

[11] Results for the euro-yen exchange rate are not reported, as price discovery in euro-dollar and dollar-yen drives this cross-rate.

4 Price Discovery: Who are the "informed" traders, humans or computers?

4.1 A return-order flow VAR

We next turn to evaluate to what extent humans or computers represent the "informed" traders in the market, in the sense of whose order flow has a larger impact on prices. To this end, we estimate return-order flow dynamics in a structural vector autoregressive (VAR) framework in the tradition of Hasbrouck (1991a), where returns are contemporaneously affected by order flow, but order flow is not contemporaneously affected by returns. Similar to Hasbrouck's (1996) decomposition of program and nonprogram order flow, we decompose order flow into two components: human-taker order flow $\left(OF^{(ht)} = HH + CH\right)$ and computer-taker order flow $\left(OF^{(ct)} = HC + CC\right)$. For each currency i, we thus estimate one return equation and two order flow equations. In light of Evans and Lyon (2008) findings on the impact of economics news on order flow in the foreign exchange market, we augment the structural VAR with U.S. macroeconomic news surprises as exogenous variables that affect both returns and order flow. Specifically, using one-minute price and order flow data, we estimate the following system of equations for each currency i,

$$
\begin{aligned}
r_{it} &= \alpha^{r} + \sum_{j=1}^{J} \beta_{ij}^{r} r_{it-j} + \sum_{j=0}^{J} \gamma_{ij}^{rct} OF_{it-j}^{(ct)} + \sum_{j=0}^{J} \gamma_{ij}^{rht} OF_{it-j}^{(ht)} + \sum_{k=1}^{K} \delta_{ik}^{r} S_{kt} + \varepsilon_{it}^{r}, \qquad (1) \\
OF_{it}^{(ht)} &= \alpha^{OF^{(ht)}} + \sum_{j=1}^{J} \beta_{ij}^{OF^{(ht)}} r_{it-j} + \sum_{j=1}^{J} \gamma_{ij}^{OF^{(ht,ht)}} OF_{it-j}^{(ht)} + \sum_{j=1}^{J} \gamma_{ij}^{OF^{(ht,ct)}} OF_{it-j}^{(ct)} + \sum_{k=1}^{K} \delta_{ik}^{OF^{(ht)}} S_{kt} + \varepsilon_{it}^{OF^{(ht)}}, \\
OF_{it}^{(ct)} &= \alpha^{OF^{(ct)}} + \sum_{j=1}^{J} \beta_{ij}^{OF^{(ct)}} r_{it-j} + \sum_{j=1}^{J} \gamma_{ij}^{OF^{(ct,ct)}} OF_{it-j}^{(ct)} + \sum_{j=1}^{J} \gamma_{ij}^{OF^{(ct,ht)}} OF_{it-j}^{(ht)} + \sum_{k=1}^{K} \delta_{ik}^{OF^{(ct)}} S_{kt} + \varepsilon_{it}^{OF^{(ct)}}.
\end{aligned}
$$

Here, r_{it} is the exchange rate return for currency i at time t; $OF_{it}^{(ht)}$ is the currency i human order flow at time t; $OF_{it}^{(ct)}$ is the currency i computer order flow at time t; and S_{kt} is the macroeconomic news announcement surprise for announcement k at time t, defined as the difference between the announcement realization and its corresponding market expectation. We use Bloomberg's real-time (unrevised) data on the expectations and realizations of $K = 28$ U.S. macroeconomic fundamentals to calculate S_{kt}. Our list of 28 announcements is based on those in Andersen et al. (2003, 2007) and Pasquariello and Vega (2007); the list is provided in Table A1 of the Appendix.[12] Since units of measurement vary across macroeconomic variables, we standardize the

[12] Our list of U.S. macroeconomic news announcements is the same as the list of announcements in Andersen et al. (2007) and Pasquariello and Vega (2007) with the addition of three announcements: unemployment rate, core PPI and core CPI. Andersen et al. (2007) and Pasquariello and Vega (2007) use International Money Market Services (MMS) data on the expectations of U.S. macroeconomic fundamentals. In contrast, we use Bloomberg data because the MMS data are not available for releases coming after 2003. Bloomberg provides survey data similar to those MMS previously provided.

resulting data surprises by dividing each of them by their sample standard deviation.[13] We estimate the VAR using minute-by-minute price and order flow data over the entire $2006-2007$ sample, as well as over the subsample of September, October, and November 2007. In both samples, 20 lags are included in the estimated VARs, i.e., $J = 20$.

The VAR specification in equation (1) does not allow human-taker order flow to contemporaneously affect computer-taker order flow or vice-versa. The advantage of this approach is that we can estimate the impulse response functions without giving more importance to a particular type of order flow, i.e., we do not need to assume a particular ordering of the human-taker and computer-taker order flow in the VAR. The disadvantage is that the human-taker and computer-taker order flow shocks may not be orthogonal. However, in our estimation this does not appear to be a problem, as our residuals are found to be approximately orthogonal (the correlation between the human-taker and computer-taker equation residuals are -0.001, -0.1 and -0.1 for the euro-dollar, dollar-yen, and euro-yen exchange rates respectively). As a further robustness check, we also estimated the VAR with two other orderings: (i) human-taker order flow affecting computer-taker order flow contemporaneously and (ii) the opposite with computer-taker order flow affecting human-taker order flow contemporaneously. The results from these two additional specifications were very similar to those obtained from equation (1), and they show that the main results presented in this section are not sensitive to alternative identification schemes in the VAR.

Before discussing the impulse response functions and the variance decompositions, we briefly summarize the main findings from the estimated coefficients in the VAR. Focusing first on the return equation, we find that one-minute returns tend to be negatively serially correlated, with the coefficient on the first own lag ranging between -0.08 and -0.15; there is thus some evidence of mean reversion in the exchange rates at these high frequencies, which is a well-known empirical finding. Both human and computer order flows are significant predictors of returns. The price impact of the lagged order flows range from around 4 to 18 basis points per billion units of order flow (denominated in the base currency). This compares to a range of about 28 to 100 basis points for the impact of the contemporaneous order flow. In the order flow equations, we find that the first own lag in both order flow equations is always highly significant, and typically around 0.1 for all currency pairs. There is thus a sizeable first-order autocorrelation in the human and computer order flows, as we had seen previously. In contrast, the coefficients on the first order cross-lags in the order flow regressions are most often substantially smaller than the coefficient on the own lags, and they vary in signs.

Finally, it is worth pointing out that despite the highly significant estimates that are recorded in the VAR estimations, the amount of variation in the order flow and return variables that is captured by their

[13] Economic theory suggests that we should also include foreign macroeconomic news announcements in equation (1). However, previous studies have found that dollar exchange rates respond little to non-U.S. macroeconomic announcements, even at high frequencies (e.g. Andersen et al., 2003). Therefore, we expect the omitted variable bias in our specification to be small.

lagged values is very limited. The R^2 for the estimated equations with only lagged variables are typically only between three to ten percent for the order flow equations, and between one and three percent for the return equations. This can be compared to an R^2 of 20 to 30 percent when one includes contemporaneous order flow in the return equations.

4.2 Impulse Response Function and Variance Decomposition Results

As originally suggested by Hasbrouck (1991a,b), we use impulse response functions to assess the price impact of the different types of order flow, and we use variance decompositions to measure the relative importance of the variables driving foreign exchange returns. In Table 3, we show the results from the impulse response analysis based on the estimation of equation (1), using the full sample for 2006-2007 and the three-month sub-sample from 2007. The size of the shock to order flow reported in Table 3 is one billion of the base currency.[14] We show both the short-run (instantaneous) impulse responses, the long-run cumulative responses, and the difference between the two responses. The long-run statistics are calculated after 30-minutes, at which point the cumulative impulse responses have converged and can thus be interpreted as the long-run total impact of the shock. All the responses are measured in basis points.

A comparison of the response to an identically-sized shock in human-taker and computer-taker order flow shows that the immediate response of prices to human order flow is often larger than the immediate response to computer order flow. This is consistent with the observation that some algorithmic trading is used for the optimal execution of large orders at a minimum cost, with computers likely breaking up larger orders and timing trades to minimize the impact on prices. We emphasize, however, that the differences in price impact between human- and computer-initiated trades, which range from 1 to 8 basis points, are not very large in economic terms. Furthermore, for some currency pairs, we find that the result can be reversed in the long-run and/or in the three-month sub-sample. Finally, an important feature of Table 3 is that the long-run response to computer order flow is, in all cases except the later sample for dollar-yen, larger than the short-run response to computer order flow. Computer order flow, on average, is therefore *not* associated with price reversals, which could be viewed as evidence that algorithmic trading does not deteriorate the informational efficiency of prices.

In Table 4, we report the fraction of the total (long-run) variance in returns that can be attributed to innovations in human order flow and computer order flow; the variance decompositions are virtually identical in the short- and long-run and we only show the long-run results. Following Hasbrouck (1991b), we interpret this variance decomposition as a summary measure of the "informativeness" of trades, and thus, in the current

[14] The one billion base currency shock is not meant to be representative of a typical trade size, but it allows for results that are in convenient units of measurement.

context, as a comparison of the relative informativeness of human and computer order flows. We find that, in the euro-dollar and dollar-yen markets, human order flow explains much more of the total variance in returns than does computer order flow. Specifically, human order flow explains about 30 percent of the total variance in returns compared to only 4 percent explained by computer order flow. Interestingly, in the cross rate, the euro-yen exchange rate, the difference between the total variance of returns explained by human order flow and that explained by computer order flow is very small. We believe this is again related to the prevalence of triangular arbitrage trading by computers in this market, with a large proportion of algorithmic trading contributing to more efficient price discovery. It is thus not surprising that in the cross rate computers and humans appear to be equally "informed."

The relatively large explanatory power of human order flow, shown by the variance decomposition results in the euro-dollar and dollar-yen markets, stands in contrast to our findings from the impulse responses discussed above. However, the difference is not unexpected. The impulse responses, by design, are calculated based on order flow shocks that are of identical sizes for computers and humans, for a given currency pair. In contrast, when analyzing what drives price variations, one must remember that human order flow accounts for about 75 percent of total volume in the euro-dollar and dollar-yen exchange rates in the full sample period and about 65 percent of total volume in the three-month sub-sample (see Figure 2). Moreover, humans tend to initiate large trades more often than computers; as a result, as was shown in Table 1, the standard deviation of human order flow is twice as large as that of computer order flow. Therefore, perhaps a more appropriate question to ask is whether, taking into account the relative shares of trades initiated by computers and humans, computers tend to contribute "disproportionately" little to the long-run variance in returns.

To answer this question we do a back-of-the-envelope calculation. We compute the relative share of the *explained* variance that is due to computer order flow as the percent of total variation in returns explained by computer order flow divided by the percent of total variation in returns explained jointly by both human and computer order flow. For example, this relative share is about $14\% = 100 \times \frac{4.74}{34}$ in the euro-dollar market (Table 4). We can then compare this relative share to the fraction of overall trading volume that is due to computer-initiated volume, also shown in the table. In the full 2006-2007 sample for the euro-dollar and the dollar-yen currency pairs, the fractions of volume due to computer takers are 22 percent and 24 percent, respectively, almost twice as large as the fractions of the explained long-run variance due to computer order flow, 14 percent and 13 percent, respectively. In the euro-yen pair, the fractions are approximately equal. These results are fairly similar in the three-month sub-sample, although the fraction of explained variance has increased somewhat compared to the fraction of volume. Thus, in the two major currency pairs, there is evidence that computer order flow contributes less to the variation in price than one would infer from just considering the proportion of trades initiated by computers. In other words, in these markets where humans and computers coexist, price discovery is still driven disproportionately by human traders.

5 How Correlated Are Algorithmic Trades and Strategies?

One concern often cited about how algorithmic trading may impact financial markets is that the strategies embodied in the trading algorithms of computers may not be as diverse as those of human traders.[15] This, in turn, could lead to large blocks of similar trading instructions reaching the market at the same time, with potential consequences of reduced liquidity and excess volatility through large temporary price impacts. Khandani and Lo (2007, 2011) describe an episode of correlated trades by hedge funds, where a number of funds and proprietary trading desks appeared to have followed similar "quantitative" long-short equity strategies, resulting in massive losses during the second week of August 2007. Khandani and Lo put forth an "unwind hypothesis", where one or more funds, for some reason, are forced to liquidate their large portfolios, which results in large adverse price impacts and subsequent losses for other funds pursuing similar strategies. The other funds subsequently reduce, or deleverage, their portfolio holdings in the same fashion, which has further negative feedback effects.

If one looks for similar episodes in our data, August 16, 2007 in the dollar-yen market stands out. It is the day with the highest realized volatility in our sample period. On that day, the Japanese yen appreciated sharply against the U.S. dollar around 6:00 a.m. and 12:00 p.m. (NY time), as shown in Figure 6. The figure also shows, for each 30-minute interval in the day, computer-taker order flow ($HC + CC$) in the top panel and human-taker order flow ($HH + CH$) in the lower panel. The two sharp exchange rate movements mentioned happened when computers, as a group, aggressively initiated sales of dollars and purchases of yen. Computers, during these periods of sharp yen appreciation, mainly traded with humans, not with other computers. Human order flow at those times was, in contrast, quite small, even though the trading volume initiated by humans (not shown) was well above that initiated by computers: human takers were therefore selling and buying dollars in almost equal amounts. The orders initiated by computers during those time intervals were therefore far more correlated than the orders initiated by humans. After 12:00 p.m., human traders, in aggregate, began to buy dollars fairly aggressively, and the appreciation of the yen against the

[15] There is little public knowledge, and no data, about the mix of strategies used by algorithmic traders in the foreign exchange market, as traders and EBS keep what they know confidential. From conversations with market participants, we believe that about half of the algorithmic trading volume on EBS over our sample period comes from traditional foreign exchange dealing banks, with the other half coming from hedge funds and commodity trading advisors (CTAs). Hedge funds and CTAs, who access EBS under prime-brokerage arrangements, can only trade algorithmically (no keyboard trading) over our sample period. Some of the banks' computer trading is related to activity on their own customer-to-dealer platforms, to automated hedging activity, and to the optimal execution of large orders. But a sizable fraction (perhaps almost a half) is believed to be proprietrary trading using a mix of strategies similar to what hedge funds and CTAs use. These strategies include various types of high-frequency arbitrage, including across different asset markets, a number of lower-frequency statistical arbitrage strategies (including carry trades), and strategies designed to automatically react to news and data releases (believed to be still fairly rare by 2007). Overall, market participants believe that the main difference between the mix of algorithmic strategies used in the foreign exchange market and the mix used in the equity market is that optimal execution algorithms are less prevalent in foreign exchange than in equity.

dollar was partially reversed.

The August 16, 2007 episode in the dollar-yen market was widely viewed at the time as the result of a sudden unwinding of large yen carry-trade positions, with many hedge funds and banks' proprietary trading desks closing risky positions and buying yen to pay back low-interest loans.[16] This is, of course, only one episode in our two-year sample, and by far the most extreme as to its impact on volatility, so one should not draw conclusions about the overall correlation of algorithmic strategies based on this single instance. Furthermore, episodes of very sharp appreciation of the yen due to the rapid unwinding of yen carry trades have occurred on a number of occasions since the late 1990s, well before algorithmic trading was allowed in this market.[17] We therefore investigate next whether there is evidence that, on average over our sample, the strategies used by computer traders have tended to be more correlated (less diverse) than those used by human traders.

5.1 Inferring correlation from trade data

We do not observe the trading strategies of market participants. However, we can infer the correlation of algorithmic strategies from the trading activity of computers and humans. The idea is the following. Traders who follow similar trading strategies and therefore send similar trading instructions at the same time, will trade less with each other than those who follow less correlated strategies. Therefore, the extent to which computers trade with each other contains information about how correlated the algorithmic strategies are.

More precisely, we consider a simple benchmark model that assumes random and independent matching of traders. This is a reasonable assumption given the lack of discrimination between keyboard traders and algorithmic traders in the EBS matching process; that is, EBS does not differentiate in any way between humans and computers when matching buy and sell orders in its electronic order book. The model allows us to determine the theoretical probabilities of the four possible trades: Human-maker/human-taker, computer-maker/human-taker, human-maker/computer-taker and computer-maker/computer-taker. We then compare these theoretical probabilities to those observed in the actual data. The benchmark model is fully described in Appendix A3, and below we merely outline the main concepts and empirical results.

Under our random and independent matching assumption, computers and humans, both of which are indifferent *ex-ante* between making and taking, trade with each other in proportion to their relative presence in the market. In a world with more human trading activity than computer trading activity (which is the

[16] A traditional carry-trade strategy borrows in a low-interest rate currency and invests in a high-interest rate currency, with the implicit assumption that the interest rate differential will not be (fully) offset by changes in the exchange rate. That is, carry trades bet on uncovered interest rate parity not holding. Although the August 16, 2007 episode occurs only a week after the events described in Khandani and Lo (2007, 2011), we are not aware of any direct link between the quant equity crisis and the carry trade unwinding.

[17] The sharp move of the yen in October 1998, which included a 1-day appreciation of the yen against the dollar of more than 7 percent, is the best-known example of the impact of the rapid unwinding of carry trades.

case in our sample), we should observe that computers take more liquidity from humans than from other computers. That is, the probability of observing human-maker/computer-taker trades, $Prob(HC)$, should be larger than the probability of observing computer-maker/computer taker trades, $Prob(CC)$. We label the ratio of the two, $Prob(HC)/\,Prob(CC)$, the computer-taker ratio, RC. Similarly, one expects humans to take more liquidity from other humans than from computers, i.e., $Prob(HH)$ should be larger than $Prob(CH)$. We label this ratio, $Prob(HH)/\,Prob(CH)$, the human-taker ratio, RH. In summary, one thus expects that $RC > 1$ and $RH > 1$, because there is more human trading activity than computer trading activity.

Importantly, the model predicts that the ratio of these two ratios, the computer-taker ratio divided by the human-taker ratio, should be equal to one. That is, the model predicts $R = RC/RH = 1$ because humans take liquidity from other humans in the same *proportion* that computers take liquidity from humans. Observing a ratio $R = RC/RH > 1$ in the data indicates that computers are trading less among themselves and more with humans than what our benchmark model predicts.[18] Therefore a ratio bigger than one can be viewed as evidence that computers have trading strategies that are more correlated than those of humans.

The test outlined above implicitly takes into account trading direction because the matching process in EBS takes it into account. Nevertheless, we also describe in detail in Appendix A3 a model that explicitly takes trading direction into account. Using notation similar to the model without trading direction, this model yields four ratios, RC^B, RC^S, RH^B, and RH^S, a computer-taker ratio where computers are buying, a computer-taker ratio where computers are selling, a human-taker ratio where humans are buying, and a human-taker ratio where humans are selling. As before, the model predicts that each of these four ratios will be greater than one, but that the ratio of the buy ratios, $R^B \equiv \frac{RC^B}{RH^B}$, and the ratio of the sell ratios, $R^S \equiv \frac{RC^S}{RH^S}$, will both be equal to one.

Based on the model described above, we calculate for each trading day in our sample a realized value for R, R^S, and R^B. Specifically, the daily realized values of RH and RC are given by $\widehat{RH} = \frac{Vol(HH)}{Vol(CH)}$ and $\widehat{RC} = \frac{Vol(HC)}{Vol(CC)}$, where, for instance, $Vol\,(HC)$ is the daily trading volume between human makers and computer takers, following the notation described in Section 2. Similarly, we define $\widehat{RH^S} = \frac{Vol\left(HH^S\right)}{Vol(CH^S)}$, $\widehat{RH^B} = \frac{Vol\left(HH^B\right)}{Vol(CH^B)}$, $\widehat{RC^S} = \frac{Vol\left(HC^S\right)}{Vol(CC^S)}$, and $\widehat{RC^B} = \frac{Vol\left(HC^B\right)}{Vol(CC^B)}$, where $Vol\left(HH^B\right)$ is the daily buy volume between human makers and human takers (i.e., buying of the base currency by the taker), $Vol\left(HH^S\right)$ is the daily sell volume between human makers and human takers, and so forth.

Table 5 shows the means of the daily ratios of ratios, $\widehat{R} = \frac{\widehat{RC}}{\widehat{RH}}$, $\widehat{R^S} = \frac{\widehat{RC^S}}{\widehat{RH^S}}$, and $\widehat{R^B} = \frac{\widehat{RC^B}}{\widehat{RH^B}}$, for each currency pair for the 2006-2007 sample, and for the three-month sub-sample. In contrast to the benchmark

[18] For the ratio R to be larger than one, either the computer-taker ratio is larger than what the model predicts or the human-taker ratio is smaller than the model predicts. For the computer-taker ratio to be larger than what the model predicts, either computers are taking too much liquidity from humans, or computers are not taking enough liquidity from other computers. Similarly, if the human-taker ratio is smaller than what the model predicts, computers are trading more with humans than the model predicts.

predictions that $R \equiv 1$, $R^B \equiv 1$ and $R^S \equiv 1$, we find that, for all three currency pairs, $\widehat{R}$, $\widehat{R^B}$ and $\widehat{R^S}$ are substantially and significantly greater than one in both sample periods. The table also shows the number of days in which the statistics are below one. In euro-dollar, $\widehat{R}$ and $\widehat{R^S}$ are above one every single day and $\widehat{R^B}$ is above one on all but one day. For dollar-yen and euro-yen, only a small fraction of the daily observations are below one. Overall, the observed ratios are highest in the cross-rate, the euro-yen, a bit lower in euro-dollar, and lowest, but still well above one, in dollar-yen. The results thus clearly show that computers do not trade with each other as much as random matching would predict. We take this as evidence that the algorithmic trading strategies used by computers are less diverse than the trading strategies used by human traders.[19]

Although this finding may raise some concerns about the impact of algorithmic trading on the foreign exchange market, a high correlation of algorithmic strategies need not necessarily be detrimental to market quality. For instance, as noted above, the evidence for a high correlation of algorithmic strategies is strongest for the euro-yen currency pair. This is consistent with a large fraction of algorithmic strategies in that currency pair being used to detect and exploit triangular arbitrage opportunities. Faced with the same price data at a particular moment, the various computers seeking to profit from the same arbitrage opportunities would certainly take the same side of the market. However, rather than create adverse effects, this would likely contribute to a more efficient price discovery process in the euro-yen market. In contrast, if the high correlation of strategies reflects a large number of algorithmic traders using the same carry trade or momentum strategies, as in the August 2007 example shown at the beginning of this section, then there may be reasons for concern. In the next section, we investigate whether there is evidence that algorithmic trading has any causal effect on market liquidity and excess volatility.

6 Causal impact of algorithmic trading on market liquidity and volatility

In the final part of our study, we attempt to explicitly identify whether the share of algorithmic trading in the foreign exchange market has a significant *causal* impact on market quality. In particular, we estimate the impact of algorithmic trading on market liquidity (depth) and volatility.

The main challenge in identifying a causal relationship between algorithmic trading (AT) and market depth and volatility is the likely reverse causality of AT and these other variables. For example, Foucault, Kadan, and Kandel's (2009) model emphasizes this endogeneity problem. In their model, algorithmic traders consume or provide liquidity depending on the level of market liquidity. At the same time, algorithmic trading

[19] We also conducted the same tests using statistics based on the number of trades of each type (HC, for instance) rather than trading volume of each type. The results were qualitatively identical, with the test statistics all statistically significantly greater than one.

activity affects market liquidity. Empirically, the reverse causality would lead to biased OLS regression coefficients, and one could not easily determine in which direction the bias would go in a regression of market depth or volatility on the fraction of algorithmic trading. To mitigate the reverse causality problem, we adopt an instrumental variable (IV) approach and we also consider Granger causality tests. Granger causality, by definition, captures causal effects in a predictive sense whereas the IV analysis focuses on the contemporaneous impact. The Granger causality tests can be performed at a minute-by-minute frequency and therefore provide a useful complement to the IV analysis, which is performed at a daily frequency, as described in detail below.

6.1 Market liquidity Granger tests

We estimate vector autoregressions using minute-by-minute data, with 20 lags, for the share of algorithmic trading activity and market depth. The share of algorithmic trading activity is measured as computer participation volume (i.e., $Vol\,(HC + CH + CC)$ as a fraction of total volume). Market depth is calculated as the total volume available to trade within five basis points on both sides of the mid-quote.[20] The measure of market depth used here is similar to the measures used by Hendershott and Riordan (2009) and Foucault and Menkveld (2008). We also include a daily linear time trend in the vector autoregressions to control for secular trends in the data, and 48 dummy variables for each 30-minute interval within the day to control for intra-day seasonality. To test for Granger causality, we consider two different test statistics. The first one is the standard Granger causality test, which is simply the F-test of whether the coefficients on all lags of the causing variable are jointly equal to zero. The second test evaluates whether the *sum* of the coefficients on the lags of the causing variable is equal to zero. Since the sum of the coefficients on the lags of the causing variable is proportional to the long-run impact of that variable, the second test can be viewed as a long-run Granger causality test. Importantly, the sum of the coefficients also tells us of the estimated direction of the (long-run) relationship. The empirical results are shown in Table 6. Each sub-panel reports the sum of the coefficients on the lags of the causing variable, as well as the F-statistics and corresponding p-values for the long-run Granger test and the traditional Granger test. Since the long-run test is associated with a clear direction in the causation, we focus on this test unless otherwise noted. The results are based on the full 2006-2007 sample.

The left panel of Table 6 indicates that, at a one-minute frequency, an increase in the share of computer activity Granger causes higher market depth in the euro-dollar and dollar-yen exchange rates. In the right panel, we observe that a scarcity of liquidity (lower depth) Granger causes a higher share of algorithmic

[20] We use this measure of depth instead of the inside bid-ask spread because the bid-ask spread in major currency pairs changes very little over the day. It is usally of the order of one or two basis points, and is a poor indicator of liquidity. Depth, by contrast, changes substantially with market conditions.

activity. These results support Foucault, Kadan, and Kandel's (2009) prediction that investors with lower monitoring costs, i.e., algorithmic traders, are likely to provide liquidity when it is scarce and are also consistent with Hendershott and Riordan's (2009) results for algorithmic trading on the Deutsche Bœrse. In the euro-yen currency pair, an increase in the share of computer activity Granger causes *lower* market depth, although this effect is not statistically significant. This may not be surprising as Figure 2 shows that computers tend to take liquidity out of the euro-yen market more often than they provide liquidity. We next turn to the IV estimation.

6.2 Market liquidity IV tests

We are interested in estimating the following regression equation,

$$Depth_{it} = \alpha_i+\beta_i AT_{it}+\gamma_i'\tau_{it}+\delta_{i,Day}Depth_{it-1}+\delta_{i,Week}\frac{1}{5}\sum_{k=1}^{5}Depth_{it-k}+\delta_{i,Month}\frac{1}{20}\sum_{k=1}^{20}Depth_{it-k}+\epsilon_{it}, \quad (2)$$

where $i = 1, 2, 3$ represents currency pairs and $t = 1, ..., T$, represents time in days. $Depth_{it}$ is the market depth as defined above, AT_{it} is the share of algorithmic trading at time t in currency pair i, τ_{it} is either a time trend or a set of time dummies that control for secular trends in the data, and ϵ_{it} is an error term. The lagged "Day", "Week", and "Month" variables are designed to control for long serial correlation in a parsimonious way, following the work of Andersen, Bollerslev, and Diebold (2007). That is, the lag structure spans 20 days—approximately one trading month—but only three coefficients need to be estimated. Since AT_{it} and ϵ_{it} may be correlated due to the potential endogeneity and reverse causality discussed above, the OLS estimator of β_i may be biased. In order to obtain an unbiased estimate, we therefore adopt an instrumental variable approach. Formally, we need to find a variable, or set of variables, z_{it}, that is uncorrelated with ϵ_{it} (validity of the instrument) and correlated with AT_{it} (relevance of the instrument).

The instrument we propose to use is the fraction of trading floors equipped to trade algorithmically on EBS relative to the total number of trading floors linked to the EBS system.[21] In order to place algorithmic trades on EBS, a special user interface is required, and the total number of trading floors with such user interfaces thus provides a measure of the overall algorithmic trading "capacity" in the market. The ratio of these algorithmic trading floors to the total number of trading floors provides a measure of the potential share of algorithmic trading. Over our sample period, it takes at least two months from the time a trader asks for the ability to trade algorithmically until algorithmic trading operation can begin. The number of trading floors of each type is therefore clearly exogenous with regards to daily depth, and the fraction of

[21] More precisely, we use a time series of the number of EBS "deal codes" of each type over our sample period. Generally speaking, EBS assigns a deal code to each trading floor equipped with at least one of its terminals, and records whether they are equipped to trade algorithmically or not. These data are confidential.

algorithmic trading floors is thus a valid instrument. In addition, it is positively correlated with the share of algorithmic trading, and it provides a relevant instrument as seen from the tests for weak instruments discussed below. Under the breakdown provided by EBS, there are three types of trading floors linked to the EBS system: purely algorithmic trading floors, purely manual trading floors, and dual trading floors, those equipped to handle both manual and algorithmic trades. We consider two natural instrumental variables, which we use simultaneously: the fraction of pure algorithmic trading floors over the total number of trading floors (including pure algorithmic, manual, and dual ones), and the fraction of the sum of pure algorithmic *and* dual trading floors over the total number of trading floors.

The data on the different types of trading floors are provided at a monthly frequency; they are transformed into daily data by repeating the monthly value each day of the month. Together with data on depth and algorithmic trading, also sampled at a daily frequency, this leads to a dataset of two years of daily data. The regression coefficient on the share of algorithmic trading will be identified from monthly variations in the instrumental variables. However, transforming the instruments to a daily frequency is more efficient than transforming all data to a monthly frequency, since the daily data help to identify the monthly shifts. In order to control for the secular trends in the data, we include either a "linear quarterly" time trend or a full set of year-quarter dummies, one for each year-quarter pair in the data (8 dummies). The linear quarterly time trend stays constant within each quarter and increases by the same amount each quarter, whereas the year-quarter dummies allows for a more flexible trend specification that can shift in arbitrary fashion from year-quarter to year-quarter. Both secular trend specifications are thus fixed within each quarter. This restriction is imposed in order to preserve the identification coming from the monthly instrumental variables. Using monthly, or finer, time dummies would eliminate the variation in the instrument and render the model unidentified. Although it is theoretically possible to include a monthly time trend, this would lead to very weak identification empirically.

The instrumental variable regressions are estimated using Limited Information Maximum Likelihood (LIML); we use LIML rather than two-stage least squares since Stock and Yogo (2005) show that the former is much less sensitive to weak instruments than the latter (see also Stock et al., 2002). We test for weak instruments by comparing the first stage F-statistic for the excluded instruments to the critical values of Stock and Yogo's (2005) test of weak instruments. These critical values are designed such that they indicate a maximal actual size for a nominal-sized 5 percent test on the coefficient; a value greater than 8.68 (5.33) for this F-statistic indicates that the maximal size of a nominal 5 percent test will be no greater than 10 (15) percent.

The regression results are presented in Table 7. The OLS results, which are likely to be biased due to the aforementioned endogeneity issues, suggest a strongly significant negative correlation between market depth and algorithmic participation. The R^2s are fairly large, reflecting the strong serial correlation in depth which is picked up by the lagged regressors. There are no systematic differences between the quarterly trend and quarterly dummies specifications.

Turning to the more interesting IV results, which control for the endogeneity bias, the coefficient estimates change fairly dramatically. The LIML-IV estimates are positive in the euro-dollar and dollar-yen market, such that increased algorithmic participation causes an increase in market depth, and negative in the euro-yen market, such that increased algorithmic participation causes a decrease in market depth. None of the IV coefficients are statistically significant, however, but it is still interesting to note that the causal relationships implied by the IV point estimates are consistent with the model of Foucault et al. (2009) and that they agree with the Granger causality results.[22] We note that even if these causal relationships were *statistically* significant, the *economic* magnitude still appears to be small.[23] Taken together, we view the IV and Granger causality results as providing some evidence that algorithmic trading slightly improves market depth in the euro-dollar and dollar-yen exchange rate markets.

6.3 Volatility and Algorithmic Trading

Apart from the effects of algorithmic trading on liquidity, an often-voiced concern is that algorithmic trading will impact price volatility. There are plausible scenarios under which algorithmic trading might lead to increased excess exchange rate volatility. For instance, the SEC/CFTC joint report on the May 2010 "flash crash" discusses the fact that a number of algorithmic traders used similar programs, leading them to take the same side of the market at the same time and accentuating market movements. Khandani and Lo (2007, 2011) also make a similar point in their papers, which highlight the high correlation of the quants' strategies in the equity market (not necessarily executed algorithmically). From a theoretical viewpoint, there may also be excess volatility if, for example, algorithmic traders, as a group, behave like the chartists described in Froot, Scharfstein, and Stein (1992) or the positive-feedback traders of DeLong, Shleifer, Summers, and Waldman (1990).[24]

Accurately measuring "excess" volatility is difficult. Here, we estimate excess volatility as volatility that cannot be explained by past volatility and the release of macroeconomic news announcements. Similar to our

[22] As seen from the first stage F-statistics for the excluded instruments, there are no signs of weak instruments, so the inference based on the IV estimates should be unbiased.

[23] Suppose that the coefficient on computer participation is about 1.5, which is in line with the IV coefficient estimates for the euro-dollar. The average monthly shift in computer participation in the euro-dollar is about 1.5 percentage points and the average market depth in the euro-dollar is about 200 million dollars. Increasing the computer participation fraction by 1.5 percentage points increases market depth by only 2.25 million dollars. Our analysis is designed to identify the effects of such monthly shifts in algorithmic trading and can therefore not be reliably used to calculate the effect of going from no algorithmic trading to a large fraction of algorithmic trading.

[24] More generally, "noise" traders as described by Bloomfield, O'Hara, and Saar (2009) cause excess volatility. Bloomfield et al. (2009) distinguish between "noise" traders and "liquidity" traders, terms that have often been used inter-changeably in the literature. "Liquidity" traders may not cause excess volatility, but "noise" traders may, according to these models.

depth analysis, we estimate vector autoregressions using minute-by-minute data, with 20 lags, for the share of algorithmic trading activity and volatility, measured as the log of the squared minute-by-minute return. In addition to the time trend and 30-minute interval dummy variables, we also use as exogenous variables the 28 macroeconomic news announcements listed in Table A1 in the Appendix. To promote tractability while simultaneously maintaining flexibility we follow Andersen et al. (2003) and impose a polynomial structure on the response patterns associated with macroeconomic news announcements. We enforce the requirement that the impact effect slowly fades to zero in the hour after the announcement. The polynomial specification is $p(\tau) = c_0[1-(\tau/60)^3] + c_1\tau[1-(\tau/60)^2] + c_2\tau^2[1-(\tau/60)]$ where τ equals 0 at the time of the announcement and τ equals 60 at sixty one minutes after the announcement. For each of the 28 announcements we only need to estimate three coefficients, c_0, c_1 and c_2, instead of sixty coefficients.

The empirical results are shown in Table 8. At the one-minute frequency, algorithmic trading appears to positively Granger cause volatility, with the exception of the dollar-yen exchange rate, where algorithmic trading appears to decrease volatility. However, the relationship is weak as the sum of the coefficients is not always statistically significant. The relationship is stronger the other way, volatility appears to positively Granger cause algorithmic trading. In other words, the evidence points more towards algorithmic traders being relatively more active in the market when volatility is high than to algorithmic traders causing volatility.

Finally, we estimate equation (2) using our instrumental variable, replacing market depth by volatility measured as the log of the sum of the squared one-minute returns for the day (i.e., daily (log-) realized volatility).[25] We also add 28 indicator variables. Indicator variable j is equal to one if macroeconomic news announcement j is released on day t, for $j = 1, ..., 28$. The results are shown in Table 9. There is some evidence of a causal effect of algorithmic trading on volatility, but it points fairly consistently towards a negative relationship, such that increased AT leads to lower volatility. As in the case of depth, however, all the point estimates also indicate that the *economic* significance of any potential relationship between AT and volatility is very small.[26] In any case, taken together, the Granger causality tests and the IV results show no systematic statistical evidence to back the often-voiced opinion that AT leads to excess volatility.

[25] Using prices sampled at a one-minute interval to calculate realized volatility has been shown not to lead to biased estimates in liquid markets (see, for instance, the results for highly-liquid stocks in Bandi and Russell, 2006, and for major exchange rates in Chaboud et al., 2010)

[26] Suppose that the coefficient on computer participation is about -0.01, which is in line with the IV coefficient estimates for the euro-dollar. The average monthly shift in computer participation in the euro-dollar is about 1.5 percentage points and the average log-volatility in the euro-dollar is about 3.76 (with returns calculated in basis points), which implies an annualized volatility of about 6.82 percent. Increasing the computer participation fraction by 1.5 percentage points decreases log-volatility by 0.015 and results in an annualized volatility of about 6.72. Thus, a typical change in computer participation might change volatility by about a tenth of a percentage point in annualized terms.

7 Conclusion

Using highly-detailed high-frequency trading data for three major exchange rates over 2006 and 2007, we analyze the impact of the growth of algorithmic trading on the spot interdealer foreign exchange market. Algorithmic trading confers a natural speed advantage over human trading, but it also limits the scope of possible trading strategies since any algorithmic strategy must be completely rule-based and pre-programmed. Our results highlight both of these features of algorithmic trading. We show that the rise of algorithmic trading in the foreign exchange market has coincided with a decrease in triangular arbitrage opportunities and with a faster response to macroeconomic news announcements. Both findings are consistent with computers having an enhanced ability to monitor and respond almost instantly to changes in the market. However, our analysis also suggests that the constraint of designing fully systematic (i.e., algorithmic) trading systems leads to less diverse strategies than otherwise, as algorithmic trades (and in the extension, strategies) are found to be more correlated than human ones.

In addition, and partly motivated by the finding that algorithmic trades tend to be more correlated than human trades, we also analyze whether the fraction of algorithmic trading in the market has any relevant causal effect on market liquidity or excess volatility. We find some evidence that algorithmic trading has a slight positive effect on market liquidity in the euro-dollar and dollar-yen market. Furthermore, we find no evidence to back the often-voiced concern that algorithmic trading leads to excessive volatility in any of the three currency pairs we analyze. One caveat we would like to add to this analysis is that our study does not cover a truly tumultuous period in financial markets; we are thus still uncertain about how algorithmic traders may behave in a crisis period.

Our price discovery analysis within a VAR framework indicates that in the euro-yen market, where speed is very important in order to capitalize on short-lived triangular arbitrage opportunities and the strategies can easily be rule-based, human and algorithmic traders appear to contribute about evenly to price discovery. In contrast, in the euro-dollar and dollar-yen markets, the two most-traded currency pairs, human traders still appear to have a dominant impact on price discovery.

Appendix

A1 Definition of Order Flow and Volume

The transactions data are broken down into categories specifying the "maker" and "taker" of the trades (human or computer), and the direction of the trades (buy or sell the base currency), for a total of eight different combinations. That is, the first transaction category may specify, say, the minute-by-minute volume of trade that results from a computer taker buying the base currency by "hitting" a quote posted by a human maker. We would record this activity as the human-computer buy volume, with the aggressor (taker) of the trade buying the base currency. The human-computer sell volume is defined analogously, as are the other six buy and sell volumes that arise from the remaining combinations of computers and humans acting as makers and takers.

From these eight types of buy and sell volumes, we can construct, for each minute, trading volume and order flow measures for each of the four possible pairs of human and computer makers and takers: human-maker/human-taker (HH), computer-maker/human-taker (CH), human-maker/computer-taker (HC), and computer-maker/computer-taker (CC). The *sum* of the buy and sell volumes for each pair gives the volume of trade attributable to that particular combination of maker and taker (denoted as $Vol(HH)$ or $Vol(HC)$, for example). The *difference* between the buy and sell volume for each pair gives the order flow attributable to that maker-taker combination (denoted as $OF(HH)$ or $OF(HC)$, for example). The sum of the four volumes, $Vol(HH + CH + HC + CC)$, gives the total volume of trade in the market. The sum of the four order flows, $OF(HH) + OF(CH) + OF(HC) + OF(CC)$, gives the total (market-wide) order flow.[27]

Throughout the paper, we use the expression "volume" and "order flow" to refer both to the market-wide volume and order flow and to the volume and order flows from other possible decompositions, with the distinction clearly indicated. Importantly, the data allow us to consider volume and order flow broken down by the type of trader who initiated the trade, human-taker ($HH + CH$) and computer-taker ($HC + CC$); by the type of trader who provided liquidity, human-maker ($HH + HC$) and computer-maker ($CH + CC$); and by whether there was any computer participation ($HC + CH + CC$).

A2 Definition of small and large trades

As discussed in the data section, we observe volume and order flow minute-by-minute rather than transaction-by-transaction. A minute is defined as a "Computer Small Buy" (CSB) if the following three conditions are

[27] There is a very high correlation in this market between trading volume per unit of time and the number of transactions per unit of time, and the ratio between the two does not vary much over our sample. Order flow measures based on amounts transacted and those based on number of trades are therefore very similar.

met: (i) computer-taker volume is larger than human-taker volume, (ii) computer-taker order flow is strictly positive, and (iii) the average computer-taker trade size in that minute is below the 75th percentile of all trades in our sample. Similarly, a minute is defined as a "Computer Large Buy" (CLB) if the first two conditions above are met and the average computer-taker trade size that minute is above the 75th percentile of trade sizes. Thus, "large" trades are specified to be clearly of above-average size. The remaining trade categories are defined in an analogous manner.

Several alternative specifications were also considered, including, for instance, amending the first condition as follows: if the share of computer-taker volume in a given minute is larger than the average share of computer-taker volume in the three months surrounding that minute, we label the minute as a computer-taker minute. Otherwise we classify it as a human-taker minute. Our results are very robust to the different specifications we tested and also robust to testing only on the most liquid hours of the day. In particular, the following results were always obtained: (a) the diagonal effect, (b) the likelihood ratio of a large trade size following a large trade size is bigger than the likelihood ratio of a small trade size following a small trade size, and (c) the likelihood ratio of computer trades following computer trades is larger than for human trades following human trades.

A3 How Correlated Are Algorithmic Trades and Strategies?

In the benchmark model, there are H_m potential human-makers (the number of humans that are standing ready to provide liquidity), H_t potential human-takers, C_m potential computer-makers, and C_t potential computer-takers. For a given period of time, the probability of a computer providing liquidity to a trader is equal to $Prob(computer - make) = \frac{C_m}{C_m+H_m}$, which we label for simplicity as α_m, and the probability of a computer taking liquidity from the market is $Prob(computer - take) = \frac{C_t}{C_t+H_t} = \alpha_t$. The remaining makers and takers are humans, in proportions $(1-\alpha_m)$ and $(1-\alpha_t)$, respectively. Assuming that these events are independent, the probabilities of the four possible trades, human-maker/human-taker, computer-maker/human-taker, human-maker/computer-taker and computer-maker/computer taker, are:

$$\begin{aligned}
Prob(HH) &= (1-\alpha_m)(1-\alpha_t) \\
Prob(HC) &= (1-\alpha_m)\alpha_t \\
Prob(CH) &= \alpha_m(1-\alpha_t) \\
Prob(CC) &= \alpha_m\alpha_t.
\end{aligned}$$

These probabilities yield the following identity,

$$Prob(HH) \times Prob(CC) \equiv Prob(HC) \times Prob(CH),$$

which can be re-written as,

$$\frac{Prob(HH)}{Prob(CH)} \equiv \frac{Prob(HC)}{Prob(CC)}.$$

We label the first ratio, $RH \equiv \frac{Prob(HH)}{Prob(CH)}$, the "human-taker" ratio and the second ratio, $RC \equiv \frac{Prob(HC)}{Prob(CC)}$, the "computer-taker" ratio. In a world with more human traders (both makers and takers) than computer traders, each of these ratios will be greater than one, because $Prob(HH) > Prob(CH)$ and $Prob(HC) > Prob(CC)$; i.e., computers take liquidity more from humans than from other computers, and humans take liquidity more from humans than from computers. However, under the baseline assumptions of our random-matching model, the identity shown above states that the ratio of ratios, $R \equiv \frac{RC}{RH}$, will be equal to one. In other words, humans will take liquidity from other humans in a similar proportion that computers take liquidity from humans.

Turning to the data, under the assumption that potential human-takers are randomly matched with potential human-makers, i.e., that the probability of a human-maker/human-taker trade is equal to the one predicted by our model, $Prob(HH) = \frac{H_m \times H_t}{(H_m+C_m)\times(H_t+C_t)}$, we can now derive implications from observations of R, our ratio of ratios. In particular, finding $R > 1$ must imply that algorithmic strategies are more correlated than what our random matching model implies. In other words, for $R > 1$ we must observe that either computers trade with each other less than expected ($Prob(CC) < \frac{C_m \times C_t}{(H_m+C_m)\times(H_t+C_t)}$) or that computers trade with humans more than expected (either $Prob(CH) > \frac{C_m \times H_t}{(H_m+C_m)\times(H_t+C_t)}$ or $Prob(HC) > \frac{H_m \times C_t}{(H_m+C_m)\times(H_t+C_t)}$).

To explicitly take into account the sign of trades, we amend the benchmark model as follows: we assume that the probability of the taker buying an asset is α_B and the probability of the taker selling is $1 - \alpha_B$. We can then write the probability of the following eight events (assuming each event is independent):

$$
\begin{aligned}
Prob(HH^B) &= (1-\alpha_m)(1-\alpha_t)\alpha_B \\
Prob(HC^B) &= (1-\alpha_m)\alpha_t\alpha_B \\
Prob(CH^B) &= \alpha_m(1-\alpha_t)\alpha_B \\
Prob(CC^B) &= \alpha_m\alpha_t\alpha_B \\
Prob(HH^S) &= (1-\alpha_m)(1-\alpha_t)(1-\alpha_B) \\
Prob(HC^S) &= (1-\alpha_m)\alpha_t(1-\alpha_B) \\
Prob(CH^S) &= \alpha_m(1-\alpha_t)(1-\alpha_B) \\
Prob(CC^S) &= \alpha_m\alpha_t(1-\alpha_B)
\end{aligned}
$$

These probabilities yield the following identities,

$$
\begin{aligned}
Prob(HH^B) \times Prob(CC^B) &\equiv Prob(HC^B) \times Prob(CH^B) \\
(1-\alpha_m)(1-\alpha_t)\alpha_m\alpha_t\alpha_B\alpha_B &\equiv (1-\alpha_m)\alpha_t\alpha_B\alpha_m(1-\alpha_t)\alpha_B \\
&\text{and} \\
Prob(HH^S) \times Prob(CC^S) &\equiv Prob(HC^S) \times Prob(CH^S) \\
(1-\alpha_m)(1-\alpha_t)(1-\alpha_B)\alpha_m\alpha_t\alpha_B(1-\alpha_B) &\equiv (1-\alpha_m)\alpha_t\alpha_B(1-\alpha_B)\alpha_m(1-\alpha_t)\alpha_B(1-\alpha_B)
\end{aligned}
$$

which can be re-written as,

$$
\begin{aligned}
\frac{Prob(HH^B)}{Prob(CH^B)} &\equiv \frac{Prob(HC^B)}{Prob(CC^B)} \\
\frac{(1-\alpha_m)(1-\alpha_t)\alpha_B}{\alpha_m(1-\alpha_t)\alpha_B} &\equiv \frac{(1-\alpha_m)\alpha_t\alpha_B}{\alpha_m\alpha_t\alpha_B} \\
&\text{and} \\
\frac{Prob(HH^S)}{Prob(CH^S)} &\equiv \frac{Prob(HC^S)}{Prob(CC^S)} \\
\frac{(1-\alpha_m)(1-\alpha_t)(1-\alpha_B)}{\alpha_m(1-\alpha_t)(1-\alpha_B)} &\equiv \frac{(1-\alpha_m)\alpha_t(1-\alpha_B)}{\alpha_m\alpha_t(1-\alpha_B)}
\end{aligned}
$$

We label the ratios, $RH^B \equiv \frac{Prob(HH^B)}{Prob(CH^B)}$, the "human-taker-buyer" ratio, $RC^B \equiv \frac{Prob(HC^B)}{Prob(CC^B)}$, the "computer-taker-buyer" ratio, $RH^S \equiv \frac{Prob(HH^S)}{Prob(CH^S)}$, the "human-taker-seller" ratio, and $RC^S \equiv \frac{Prob(HC^S)}{Prob(CC^S)}$, the "computer-taker-seller" ratio.

In a world with more human traders (both makers and takers) than computer traders, each of these ratios will be greater than one, because $Prob(HH^B) > Prob(CH^B)$, $Prob(HH^S) > Prob(CH^S)$, $Prob(HC^B) > Prob(CC^B)$, and $Prob(HC^S) > Prob(CC^S)$. That is, computers take liquidity more from humans than from other computers, and humans take liquidity more from humans than from computers. However, under the baseline assumptions of our random-matching model, the identity shown above states that the ratio of ratios, $R^B \equiv \frac{RC^B}{RH^B}$, will be equal to one, and $R^S \equiv \frac{RC^S}{RH^S}$, will also be equal to one.

Under the assumption that potential human-taker-buyers are randomly matched with potential human-maker-sellers and human-taker-sellers are randomly matched with potential human-maker-buyers, i.e., that the probability of a human-maker-seller/human-taker-buyer trade is equal to the one predicted by our model, $Prob(HH^B) = (1 - \alpha_m)(1 - \alpha_t)\alpha_B$, and the probability of a human-maker-buyer/human-taker-seller trade is equal to the one predicted by our model, $Prob(HH^S) = (1 - \alpha_m)(1 - \alpha_t)(1 - \alpha_B)$, we can now derive implications from observations of $R^B$and R^S our ratio of ratios. In particular, finding $R^B > 1$ must imply that algorithmic strategies of buyers are more correlated than what our random matching model implies. In other words, for $R^B > 1$ we must observe that either computers trade with each other less than expected when they are buying ($Prob(CC^B) < \alpha_m\alpha_t\alpha_B$) or that computers trade with humans more than expected when they are buying (either $Prob(CH^B) > \alpha_m(1-\alpha_t)\alpha_B$ or $Prob(HC^B) > (1-\alpha_m)\alpha_t\alpha_B$). Symmetrically, for $R^S > 1$ we must observe that either computers trade with each other less than expected when they are selling ($Prob(CC^S) < \alpha_m\alpha_t(1 - \alpha_B)$) or that computers trade with humans more than expected when they are selling (either $Prob(CH^S) > \alpha_m(1 - \alpha_t)(1 - \alpha_B)$ or $Prob(HC^S) > (1 - \alpha_m)\alpha_t(1 - \alpha_B)$).

Table A1: Macroeconomic announcements. The table shows a list of the scheduled macroeconomic announcements that are included in the VAR regressions in equation (1).

Quarterly Announcements	
	GDP advance
	GDP preliminary
	GDP final
Monthly Announcements	
Real Activity:	
	Unemployment rate
	Nonfarm payroll employment
	Retail sales
	Industrial production
	Capacity utilization
	Personal income
	Consumer credit
Consumption:	
	New home sales
	Personal consumption expenditures
Investment:	
	Durable goods orders
	Construction spending
	Factory orders
	Business inventories
Government purchases:	
	Treasury budget
Net exports:	
	Trade balance
Prices:	
	Producer price index
	Core PPI
	Consumer price index
	Core CPI
Forward looking:	
	Consumer confidence index
	NAPM index
	Housing starts
	Index of leading indicators
Six-week announcements	
	Target federal funds rate
Weekly announcements	
	Initial unemployment claims

References

[1] Andersen, T.G., T. Bollerslev, and F.X. Diebold, 2007, Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility, *Review of Economics and Statistics* 89, 701-720.

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

[3] Andersen, T.G., T. Bollerslev, F.X. Diebold, and C. Vega, 2007. Real-time price discovery in global stock, bond, and foreign exchange markets, *Journal of International Economics* 73, 251-277.

[4] Bandi, F.M., and J.R. Russell, 2006, Separating Microstructure Noise from Volatility, *Journal of Financial Economics*, 79, 655-692.

[5] Bertsimas, D., and Lo, A., 1998, Optimal Control of Execution Costs, *Journal of Financial Markets* 1, 1–50.

[6] Biais, B., Foucault, T., and Moinas S., 2010, Equilibrium Algorithmic Trading, Working Paper, Tolouse School of Economics.

[7] Biais, B., Hillion, P., and Spatt, C., 1995, An empirical analysis of the limit order book and the order flow in the Paris Bourse, *Journal of Finance*, 50, 1655-1689.

[8] Bloomfield, R., O'Hara, M., and Saar, G., 2009, How noise trading affects markets: an experimental analysis, Review of Financial Studies, 22, 2275-2302.

[9] Brogaard, J. A., 2010, High Frequency Trading and Its Impact on Market Quality, working paper, Northwestern University.

[10] Chaboud, A., Chiquoine, B., Hjalmarsson, E., Loretan, M., 2010, Frequency of Observations and the Estimation of Integrated Volatility in Deep and Liquid Financial Markets, *Journal of Empirical Finance*, 17, 212-240.

[11] Chaboud, A., Chernenko, S., and Wright, J., 2008, Trading Activity and Macroeconomic Announcements in High-Frequency Exchange Rate Data, *Journal of the European Economic Association*, 6, 589-596.

[12] DeLong, B., Shleifer, A., Summers, L., Waldmann, R.J., 1990, Positive Feedback Investment Strategies and Destabilizing Rational Speculation, *Journal of Finance*, 45, 379-395.

[13] Easley, D., and O'Hara, M., 1987, Price, trade size, and information in securities markets, *Journal of Financial Economics*, 19, 69-90.

[14] Evans, M., and R. Lyons, 2008. How is Macro News Transmitted to Exchange Rates?, *Journal of Financial Economics* 88, 26-50.

[15] Foucault, T., Kadan, O., and Kandel, E., 2009, Liquidity Cycles and Make/Take Fees in Electronic Markets, working paper, Washington University in St. Louis.

[16] Foucault, T., and Menkveld, A., 2008, Competition for Order Flow and Smart Order Routing Systems, *Journal of Finance*, 63, 119-158.

[17] Froot, K., Scharfstein, D. S., Stein, J.C., 1992, Herd on the Street: Informational Inefficiencies in a Market with Short-Term Speculation, *Journal of Finance*, 47, 1461-1484.

[18] Hasbrouck, J., 1991a. Measuring the information content of stock trades, *Journal of Finance* 46, 179-207.

[19] Hasbrouck, J., 1991b. The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies* 4, 571-595.

[20] Hasbrouck, J., 1996. Order characteristics and stock price evolution: An application to program trading, *Journal of Financial Economics* 41, 129-149.

[21] He, H., and J. Wang, 1995, Differential information and dynamic behavior of stock trading volume, *Review of Financial Studies*, 8, 919-972.

[22] Hendershott, T., C.M. Jones, and A.J. Menkveld, 2011. Does algorithmic trading improve liquidity?, *Journal of Finance,* 66, 1-33.

[23] Hendershott, T., and Riordan, R., 2009, Algorithmic Trading and Information, working paper, University of California at Berkeley.

[24] Khandani, A., and Lo, A., 2007, What Happened to the Quants in August 2007?, *Journal of Investment Management* 5, 5-54.

[25] Khandani, A., and Lo, A., 2011, What Happened to the Quants in August 2007? Evidence from Factors and Transactions Data, *Journal of Financial Markets* 14, 1-46.

[26] Lo, A., Repin, D.V., and Steenbarger, B.N., 2005., Fear and Greed in Financial Markets: A Clinical Study of Day-Traders, *American Economic Review* 95, 352–359.

[27] Pasquariello, P., and Vega, C., 2007, Informed and Strategic Order Flow in the Bond Markets, *Review of Financial Studies*, 20, 1975-2019.

[28] Stock, J.H., J.H. Wright, and M. Yogo, 2002. A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments, *Journal of Business and Economic Statistics*, 20, 518-529.

[29] Stock, J.H., and M. Yogo, 2005. Testing for Weak Instruments in Linear IV Regression, in D.W.K. Andrews and J.H. Stock, eds., *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge: Cambridge University Press, 80–108.

[30] Zhang, X.F., 2010, The Effect of High-Frequency Trading on Stock Volatility and Price Discovery, working paper, Yale University.

Table 1: Summary statistics for the one-minute order flow data. The mean and standard deviation, as well as the first-order autocorrelation, ρ, are shown for each variable and currency pair. The order flows are expressed in millions of the base currency, and the summary statistics are given for both the full 2006-2007 sample, as well as for the three-month sub-sample, which only uses observations from September, October, and November of 2007. The first row for each currency pair shows the summary statistics for the total market-wide order flow. The following two rows give the results for the order flows broken down into human-takers and computer-takers and the last four rows show the results for the order flow decomposed into each maker-taker pair. There are a total of 717,120 observations in the full two-year sample and 89,280 observations in the three-month sub sample. We show the statistical significance of the first order autocorrelation. The ***, **, and * represent significance at the 1, 5, and 10 percent level, respectively.

	Full 2006-2007 Sample			3-month sub sample		
Variable	Mean	Std. dev.	ρ	Mean	Std. dev.	ρ
	USD/EUR					
Total order flow $(HH+CH+HC+CC)$	0.0315	25.95	0.150***	−0.0937	29.71	0.174***
H-taker order flow $(HH+CH)$	0.0413	23.98	0.155***	−0.0796	26.81	0.189***
C-taker order flow $(HC+CC)$	−0.0099	9.94	0.127***	−0.0140	12.89	0.115***
H-maker/H-taker order flow (HH)	0.1425	19.96	0.177***	0.0327	21.92	0.209***
C-maker/H-taker order flow (CH)	−0.1012	8.90	0.166***	−0.1123	10.76	0.215***
H-maker/C-taker order flow (HC)	0.0123	8.92	0.152***	0.0483	11.59	0.150***
C-maker/C-taker order flow (CC)	−0.0222	2.79	0.053***	−0.0623	3.95	0.072***
	JPY/USD					
Total order flow $(HH+CH+HC+CC)$	0.1061	20.10	0.189***	−0.3439	23.64	0.211***
H-taker order flow $(HH+CH)$	0.0853	19.11	0.190***	−0.2088	22.03	0.204***
C-taker order flow $(HC+CC)$	0.0209	8.39	0.170***	−0.1351	11.59	0.158***
H-maker/H-taker order flow (HH)	0.1037	16.00	0.209***	−0.1203	17.46	0.226***
C-maker/H-taker order flow (CH)	−0.0184	6.90	0.172***	−0.0885	9.18	0.162***
H-maker/C-taker order flow (HC)	0.0198	7.57	0.198***	−0.0901	10.17	0.191***
C-maker/C-taker order flow (CC)	0.0011	2.45	0.032***	−0.045	3.87	0.026***
	JPY/EUR					
Total order flow $(HH+CH+HC+CC)$	−0.0648	7.09	0.152***	−0.1574	8.60	0.147***
H-taker order flow $(HH+CH)$	−0.0497	5.70	0.150***	−0.1216	6.21	0.125***
C-taker order flow $(HC+CC)$	−0.0151	4.84	0.146***	−0.0358	6.70	0.131***
H-maker/H-taker order flow (HH)	−0.0172	4.42	0.159***	−0.0600	4.31	0.157***
C-maker/H-taker order flow (CH)	−0.0325	2.89	0.129***	−0.0616	3.72	0.092***
H-maker/C-taker order flow (HC)	−0.0095	4.53	0.173***	−0.0264	6.10	0.161***
C-maker/C-taker order flow (CC)	−0.0056	1.56	0.023***	−0.0095	2.56	−0.001

Table 2: Trade frequency conditional on previous trade. We report the conditional frequency of observing a given trade type, conditioning on the trade type in the previous period. Eight different trades types are considered: "Human Small Buy" and "Human Small Sell" orders (denoted HSB and HSS, respectively), "Human Large Buy" and "Human Large Sell" orders (denoted HLB and HLS, respectively), and the corresponding computer trades (denoted CSB, CSS, CLB, and CLS, respectively). "Z" denotes the no-trade event and the "Unconditional" row reports the average unconditional frequency of a given trade type. The "Likelihood ratio"row reports the diagonal likelihood ratio for the trade type indicated in the column header, and is defined as the ratio between the probability of observing that trade type given that the same trade type was observed in the previous period (e.g., observing HSB in period t after observing HSB in period $t-1$) and the probability of observing that trade type given that a different trade type was observed in the previous period. That is, the (diagonal) likelihood ratio for the HSB column is defined as $\Pr(HSB_t|HSB_{t-1})/\Pr(HSB_t|HSB_{t-1}^C)$, where HSB_{t-1}^C denotes the complement of HSB_{t-1}; the likelihood ratios for the other trade types are defined analogously. The largest conditional frequency in each column is indicated in bold. The results are based on the the three-month sub-sample, which only uses data from September, October, and November of 2007.

	USD/EUR								
	HSB(t)	HSS(t)	HLB(t)	HLS(t)	CSB(t)	CSS(t)	CLB(t)	CLS(t)	Z(t)
HSB(t-1)	**0.32**	0.24	0.09	0.07	0.09	0.09	0.02	0.03	0.06
HSS(t-1)	0.26	**0.29**	0.08	0.09	0.09	0.08	0.03	0.02	0.06
HLB(t-1)	0.26	0.21	**0.17**	0.12	0.07	0.07	0.03	0.03	0.05
HLS(t-1)	0.23	0.25	0.12	**0.18**	0.07	0.06	0.02	0.03	0.04
CSB(t-1)	0.24	0.22	0.08	0.06	**0.17**	0.11	0.04	0.03	0.06
CSS(t-1)	0.25	0.22	0.07	0.07	0.11	**0.15**	0.03	0.03	0.06
CLB(t-1)	0.23	0.21	0.10	0.08	0.13	0.09	**0.06**	0.04	0.05
CLS(t-1)	0.24	0.20	0.08	0.10	0.11	0.12	0.03	**0.07**	0.05
Z(t-1)	0.22	0.19	0.05	0.05	0.08	0.08	0.02	0.02	**0.30**
Unconditional	0.27	0.24	0.09	0.09	0.10	0.09	0.03	0.03	0.07
Likelihood ratio	1.28	1.30	2.00	2.23	1.85	1.86	2.56	2.60	5.48
	JPY/USD								
	HSB(t)	HSS(t)	HLB(t)	HLS(t)	CSB(t)	CSS(t)	CLB(t)	CLS(t)	Z(t)
HSB(t-1)	**0.32**	0.22	0.10	0.08	0.09	0.10	0.02	0.03	0.05
HSS(t-1)	0.25	**0.29**	0.08	0.10	0.11	0.08	0.03	0.02	0.04
HLB(t-1)	0.25	0.19	**0.17**	0.11	0.09	0.08	0.03	0.03	0.05
HLS(t-1)	0.21	0.25	0.11	**0.17**	0.09	0.08	0.03	0.03	0.04
CSB(t-1)	0.18	0.17	0.07	0.06	**0.22**	0.17	0.06	0.04	0.04
CSS(t-1)	0.19	0.18	0.06	0.07	0.17	**0.20**	0.04	0.05	0.03
CLB(t-1)	0.17	0.16	0.10	0.07	0.18	0.15	**0.08**	0.04	0.05
CLS(t-1)	0.16	0.16	0.09	0.09	0.16	0.18	0.05	**0.08**	0.04
Z(t-1)	0.19	0.15	0.08	0.07	0.08	0.07	0.03	0.02	**0.31**
Unconditional	0.24	0.22	0.09	0.09	0.12	0.12	0.03	0.03	0.06
Likelihood ratio	1.50	1.45	1.99	2.08	2.00	1.97	2.41	2.53	7.38
	JPY/EUR								
	HSB(t)	HSS(t)	HLB(t)	HLS(t)	CSB(t)	CSS(t)	CLB(t)	CLS(t)	Z(t)
HSB(t-1)	**0.15**	0.10	0.08	0.06	0.16	0.15	0.05	0.06	0.19
HSS(t-1)	0.13	**0.14**	0.06	0.07	0.18	0.14	0.06	0.05	0.18
HLB(t-1)	0.13	0.10	**0.13**	0.08	0.14	0.12	0.06	0.05	0.18
HLS(t-1)	0.12	0.12	0.07	**0.12**	0.13	0.14	0.06	0.06	0.18
CSB(t-1)	0.08	0.07	0.03	0.03	**0.31**	0.23	0.09	0.07	0.09
CSS(t-1)	0.07	0.07	0.03	0.04	0.26	**0.30**	0.07	0.08	0.08
CLB(t-1)	0.07	0.07	0.05	0.04	0.27	0.20	**0.12**	0.09	0.09
CLS(t-1)	0.07	0.07	0.03	0.05	0.23	0.24	0.10	**0.11**	0.09
Z(t-1)	0.08	0.07	0.04	0.04	0.08	0.07	0.03	0.03	**0.56**
Unconditional	0.09	0.08	0.05	0.05	0.20	0.18	0.07	0.06	0.22
Likelihood ratio	1.81	1.75	2.94	2.85	1.78	1.88	1.89	1.91	4.67

Table 3: Impulse responses from the VAR specification with human-taker and computer-taker order flow. The table shows the impulse responses for returns, in basis points, as a result of a one-billion base-currency shock to the human-taker order flow ($HH + CH$) or computer-taker ($CC + HC$) order flow, denoted H-taker and C-taker in the table headings, respectively. The results are based on estimation of equation (1), using minute-by-minute data. We show the results for the full 2006-2007 sample and for the three-month sub-sample, which only uses data from September, October, and November of 2007. For each currency pair we show the short-run (immediate) response of returns; the corresponding cumulative long-run response of returns, calculated as the cumulative impact of the shock after 30 minutes; and the difference between the cumulative long-run response in returns minus the immediate response of returns, i.e., we provide the extent of over-reaction or under-reaction to an order flow shock. There are a total of 717,120 minute-by-minute observations in the full two-year sample and 89,280 observations in the three-month sub-sample. We show in parenthesis the standard errors of the difference between the short-run and long-run response. These standard errors are calculated by bootstrapping, using 200 repetitions.

	Full 2006-2007 sample		3-month sub-sample	
	H-taker	C-taker	H-taker	C-taker
	USD/EUR			
Short run	28.06	26.94	23.20	25.22
Long run	27.87	32.35	24.16	31.38
Difference	−0.20	5.42	0.96	6.16
	(0.29)	(0.67)	(0.72)	(1.36)
	JPY/USD			
Short run	46.77	39.81	48.02	44.89
Long run	47.50	44.27	49.54	40.63
Difference	0.74	4.46	1.52	−4.26
	(0.48)	(1.08)	(1.36)	(2.35)
	JPY/EUR			
Short run	99.32	102.71	124.02	115.52
Long run	108.07	109.85	132.53	123.26
Difference	8.75	7.14	8.51	7.74
	(1.50)	(1.67)	(4.79)	(4.76)

Table 4: Variance decompositions from the VAR specification with human-taker and computer-taker order flow. The table provides the long-run variance decomposition of returns, expressed in percent and calculated at the 30 minute horizon, based on estimation of equation (1), using minute-by-minute data. That is, the table shows the proportion of the long-run variation in returns that can be attributed to shocks to the human-taker order flow ($HH + CH$) and the computer-taker ($CC + HC$) order flow, denoted H-taker and C-taker in the table headings, respectively. For each currency pair we show the actual variance decomposition, and the proportion of the *explained* variance in returns that can be attributed to each order flow type. That is, we re-scale the variance decompositions so that they add up to 100 percent. In addition, we show the average proportion of trading volume attributed to each trading type. We show results for the full 2006-2007 sample and for the three-month sub-sample, which only uses data from September, October, and November of 2007. There are a total of $717,120$ minute-by-minute observations in the full two-year sample and $89,280$ observations in the three-month sub-sample. We show in parenthesis the standard errors calculated by bootstrapping, using 200 repetitions.

	Full 2006-2007 sample		3-month sub-sample	
	H-taker	C-taker	H-taker	C-taker
	USD/EUR			
Variance decomposition	29.27	4.74	25.92	7.25
	(0.95)	(0.19)	(0.79)	(0.42)
Proportion of *explained* share	86.06	13.94	78.14	21.86
	(2.79)	(0.56)	(2.38)	(1.27)
Proportion of trading volume	78	22	69	31
	JPY/USD			
Variance decomposition	29.31	4.22	28.59	7.22
	(0.35)	(0.11)	(0.50)	(0.33)
Proportion of *explained* share	87.41	12.59	79.84	20.16
	(1.04)	(0.33)	(1.40)	(0.92)
Proportion of trading volume	76	24	65	35
	JPY/EUR			
Variance decomposition	12.03	9.28	12.47	12.67
	(0.21)	(0.20)	(0.38)	(0.38)
Proportion of *explained* share	56.45	43.55	49.60	50.40
	(0.99)	(0.94)	(1.51)	(1.51)
Proportion of trading volume	56	44	40	59

Table 5: Tests of correlation among algorithmic trading strategies. The table report estimates of the relative degree to which computers trade with each other compared to how much they trade with humans, based on the benchmark model described in the main text. In particular, we report the mean estimates of the daily ratio $R = RC/RH$, with standard errors shown in parentheses below the estimates, where $R > 1$ indicates that computers trade less with each other than random matching would predict. The table also reports mean estimates of the daily buy ratio, $R^B = RC^B/RH^B$, and the daily sell ratio, $R^S = RC^S/RH^S$, based on the model that explicitly takes into account the trading direction. Again, $R^B > 1$ and $R^S > 1$ indicate that computers trade less with each other than random matching would predict. We also show the number of days that had a ratio that was less than one. We report the results for the full 2006-2007 sample and the three-month sub-sample, which only uses data from September, October, and November of 2007. The ***, **, and * represent a statistically significant deviation from one at the 1, 5, and 10 percent level, respectively.

	Full 2006-2007 sample			3-month sub sample		
	R	R^S	R^B	R	R^S	R^B
	USD/EUR					
Mean of daily $R = RC/RH$	1.4463***	1.4526***	1.4396***	1.3721***	1.3591***	1.3843***
Standard Error	(0.0063)	(0.0074)	(0.0072)	(0.0099)	(0.0139)	(0.0122)
No. of days with $R < 1$	0	0	1	0	0	0
No. of obs	498	498	498	62	62	62
	JPY/USD					
Mean of daily $R = RC/RH$	1.2619***	1.2731***	1.2482***	1.1719***	1.1830***	1.1667***
Standard Error	(0.0074)	(0.0088)	(0.0082)	(0.0142)	(0.0176)	(0.0164)
No. of days with $R < 1$	15	29	33	4	8	3
No. of obs	498	498	498	62	62	62
	JPY/EUR					
Mean of daily $R = RC/RH$	1.6886***	1.6964***	1.6875***	1.6242***	1.6420***	1.6168***
Standard Error	(0.0154)	(0.0162)	(0.0175)	(0.0250)	(0.0291)	(0.0283)
No. of days with $R < 1$	4	4	10	0	0	0
No. of obs	498	498	498	62	62	62

Table 6: Granger causality tests for market depth. We report tests of whether algorithmic trading Granger causes market depth (left hand panel) and whether market depth Granger causes algorithmic trading (right hand panel). The tests are based on intra-daily minute-by-minute data. Algorithmic trading is measured as computer participation volume relative to total volume (AT participation). The five rows in each sub-panel presents the following results. The first row reports the sum of the lag-coefficients for the causing variable. The second and third row report the corresponding F-statistic and p-value for the null hypothesis that this sum is equal to zero. The fourth and fifth row report the F-statistic and p-value for the null hypothesis that the coefficients on all lags of the causing variable are jointly equal to zero. Data for the sample spanning 2006 and 2007 are used, with 717,120 observations in total.

	USD/EUR	JPY/USD	JPY/EUR		USD/EUR	JPY/USD	JPY/EUR
Tests of AT Participation Granger causing Depth				Tests of Depth Granger causing AT Participation			
Sum of AT coeffs.	0.029	0.046	−0.003	Sum of Depth coeffs.	−0.005	−0.007	−0.036
F-stat: Sum of AT coeffs.=0	11.950	50.785	0.415	F-stat: Sum of Depth coeffs.=0	161.170	198.657	318.110
p-value	0.001	0.000	0.520	p-value	0.000	0.000	0.000
F-stat: All AT coeffs.=0	3.415	7.561	2.832	F-stat: All Depth coeffs.=0	38.655	51.432	40.947
p-value	0.000	0.000	0.000	p-value	0.000	0.000	0.000

Table 7: Regressions of market depth on the fraction of algorithmic trading. The table shows the results from estimating the relationship between market depth and the fraction of algorithmic trading as specified in equation (2), using daily data from 2006 and 2007. In all cases, only the coefficient on the fraction of algorithmic trading is displayed and robust standard errors are given in parentheses below the coefficient estimates. The left hand side of the table shows the results with a quarterly time trend included in the regressions and the right hand side of the table shows the results with year-quarter time dummies (i.e., eight time dummies, one for each quarter in the two years of data) included in the regressions. The fraction of algorithmic trading is measured as the fraction of the total trade volume that has a computer involved on at least one side of the trade (i.e., as a maker or a taker). The top part of each panel shows the results from a standard OLS estimation, along with the R^2, and the lower part of each panel shows the results from the IV specification estimated with Limited Information Maximum Likelihood (LIML). For the IV specification, the Stock and Yogo (2005) F-test of weak instruments are also shown. The critical values for Stock and Yogo's (2005) F-test are designed such that they indicate a maximal actual size for a nominal sized five percent test on the coefficient in the LIML estimation. In order for the actual size of the LIML test to be no greater than 10% (15%), the F-statistic should exceed 8.68 (5.33) in the case considered here, with two excluded instruments and one endogenous regressor. There are a total of 498 daily observations in the data. The ***, **, and * represent significance at the 1, 5, and 10 percent level, respectively.

	With quarterly time trend			With year-quarter time dummies		
	USD/EUR	JPY/USD	JPY/EUR	USD/EUR	JPY/USD	JPY/EUR
	OLS estimation					
Coeff. on AT	-1.269^{***}	-1.274^{***}	-0.566^{***}	-1.922^{***}	-1.478^{***}	-0.826^{***}
	(0.416)	(0.310)	(0.103)	(0.488)	(0.352)	(0.133)
R^2 (%)	78.89%	56.07%	60.20%	80.22%	57.97%	63.20%
	LIML-IV estimation					
Coeff. on AT	1.625	2.154	-0.081	1.254	0.632	-0.230
	(1.186)	(1.310)	(0.244)	(1.289)	(0.953)	(0.294)
F-Stat	21.75	17.92	36.04	37.02	23.76	14.64

Table 8: Granger causality tests for exchange rate volatility. We report tests of whether algorithmic trading Granger causes exchange rate volatility (left hand panel) and whether exchange rate volatility Granger causes algorithmic trading (right hand panel). The tests are based on intra-daily minute-by-minute data. Algorithmic trading is measured as computer participation volume relative to total volume (AT participation). The five rows in each sub-panel presents the following results. The first row reports the sum of the lag-coefficients for the causing variable. The second and third row report the corresponding F-statistic and p-value for the null hypothesis that this sum is equal to zero. The fourth and fifth row report the F-statistic and p-value for the null hypothesis that the coefficients on all lags of the causing variable are jointly equal to zero. Data for the sample spanning 2006 and 2007 are used, with 717,120 observations in total.

	USD/EUR	JPY/USD	JPY/EUR		USD/EUR	JPY/USD	JPY/EUR
Tests of AT Participation Granger causing Volatility				Tests of Volatility Granger causing AT Participation			
Sum of AT coeffs.	0.008	−0.003	0.004	Sum of Volatility coeffs.	1.128	0.631	0.941
F-stat: Sum of AT coeffs.=0	11.986	3.833	9.648	F-stat: Sum of Volatility coeffs.=0	21.117	11.495	29.916
p-value	0.001	0.050	0.002	p-value	0.000	0.001	0.000
F-stat: All AT coeffs.=0	1.468	1.223	1.713	F-stat: All Volatility coeffs.=0	2.328	1.334	2.233
p-value	0.081	0.223	0.024	p-value	0.001	0.145	0.001

Table 9: Regressions of exchange rate volatility on the fraction of algorithmic trading. The table shows the results from estimating the relationship between exchange rate volatility and the fraction of algorithmic trading as specified in equation (2), using daily data from 2006 and 2007. In all cases, only the coefficient on the fraction of algorithmic trading is displayed and robust standard errors are given in parentheses below the coefficient estimates. The left hand side of the table shows the results with a quarterly time trend included in the regressions and the right hand side of the table shows the results with year-quarter time dummies (i.e., eight time dummies, one for each quarter in the two years of data) included in the regressions. The fraction of algorithmic trading is measured as the fraction of the total trade volume that has a computer involved on at least one side of the trade (i.e., as a maker or a taker). The top part of each panel shows the results from a standard OLS estimation, along with the R^2, and the lower part of each panel shows the results from the IV specification estimated with Limited Information Maximum Likelihood (LIML). For the IV specification, the Stock and Yogo (2005) F-test of weak instruments are also shown. The critical values for Stock and Yogo's (2005) F-test are designed such that they indicate a maximal actual size for a nominal sized five percent test on the coefficient in the LIML estimation. In order for the actual size of the LIML test to be no greater than 10% (15%), the F-statistic should exceed 8.68 (5.33) in the case considered here, with two excluded instruments and one endogenous regressor. There are a total of 498 daily observations in the data. The ***, **, and * represent significance at the 1, 5, and 10 percent level, respectively.

	With quarterly time trend			With year-quarter time dummies		
	USD/EUR	JPY/USD	JPY/EUR	USD/EUR	JPY/USD	JPY/EUR
	OLS estimation					
Coeff. on AT	0.0025	0.0001	0.0032^{**}	0.0072^{***}	0.0008	0.0065^{***}
	(0.0021)	(0.0020)	(0.0012)	(0.0023)	(0.0024)	(0.0017)
R^2 (%)	60.20%	62.33%	71.24%	63.38%	63.81%	72.76%
	LIML-IV estimation					
Coeff. on AT	-0.0148^{**}	-0.0203^{**}	-0.0024	-0.0111^{*}	-0.0118^{*}	-0.0118
	(0.0061)	(0.0091)	(0.0040)	(0.0060)	(0.0069)	(0.0171)
F-Stat	26.13	17.80	33.47	31.22	20.58	2.42

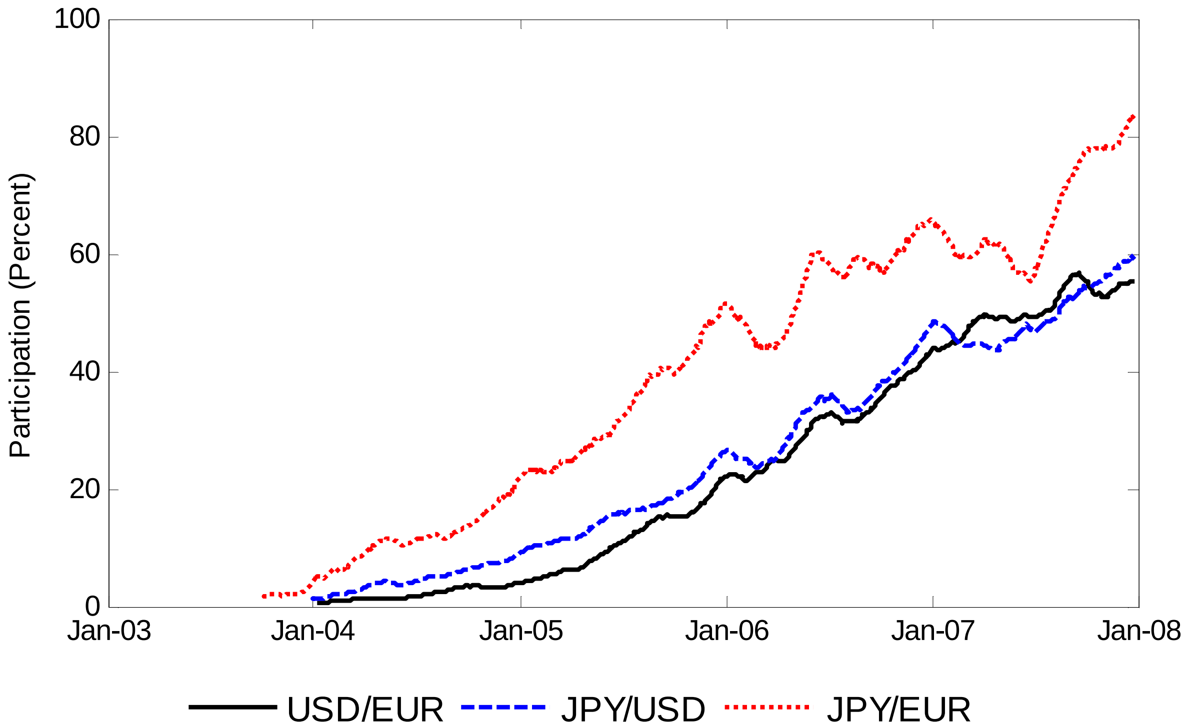

Figure 1: 50-day moving averages of participation rates of algorithmic traders

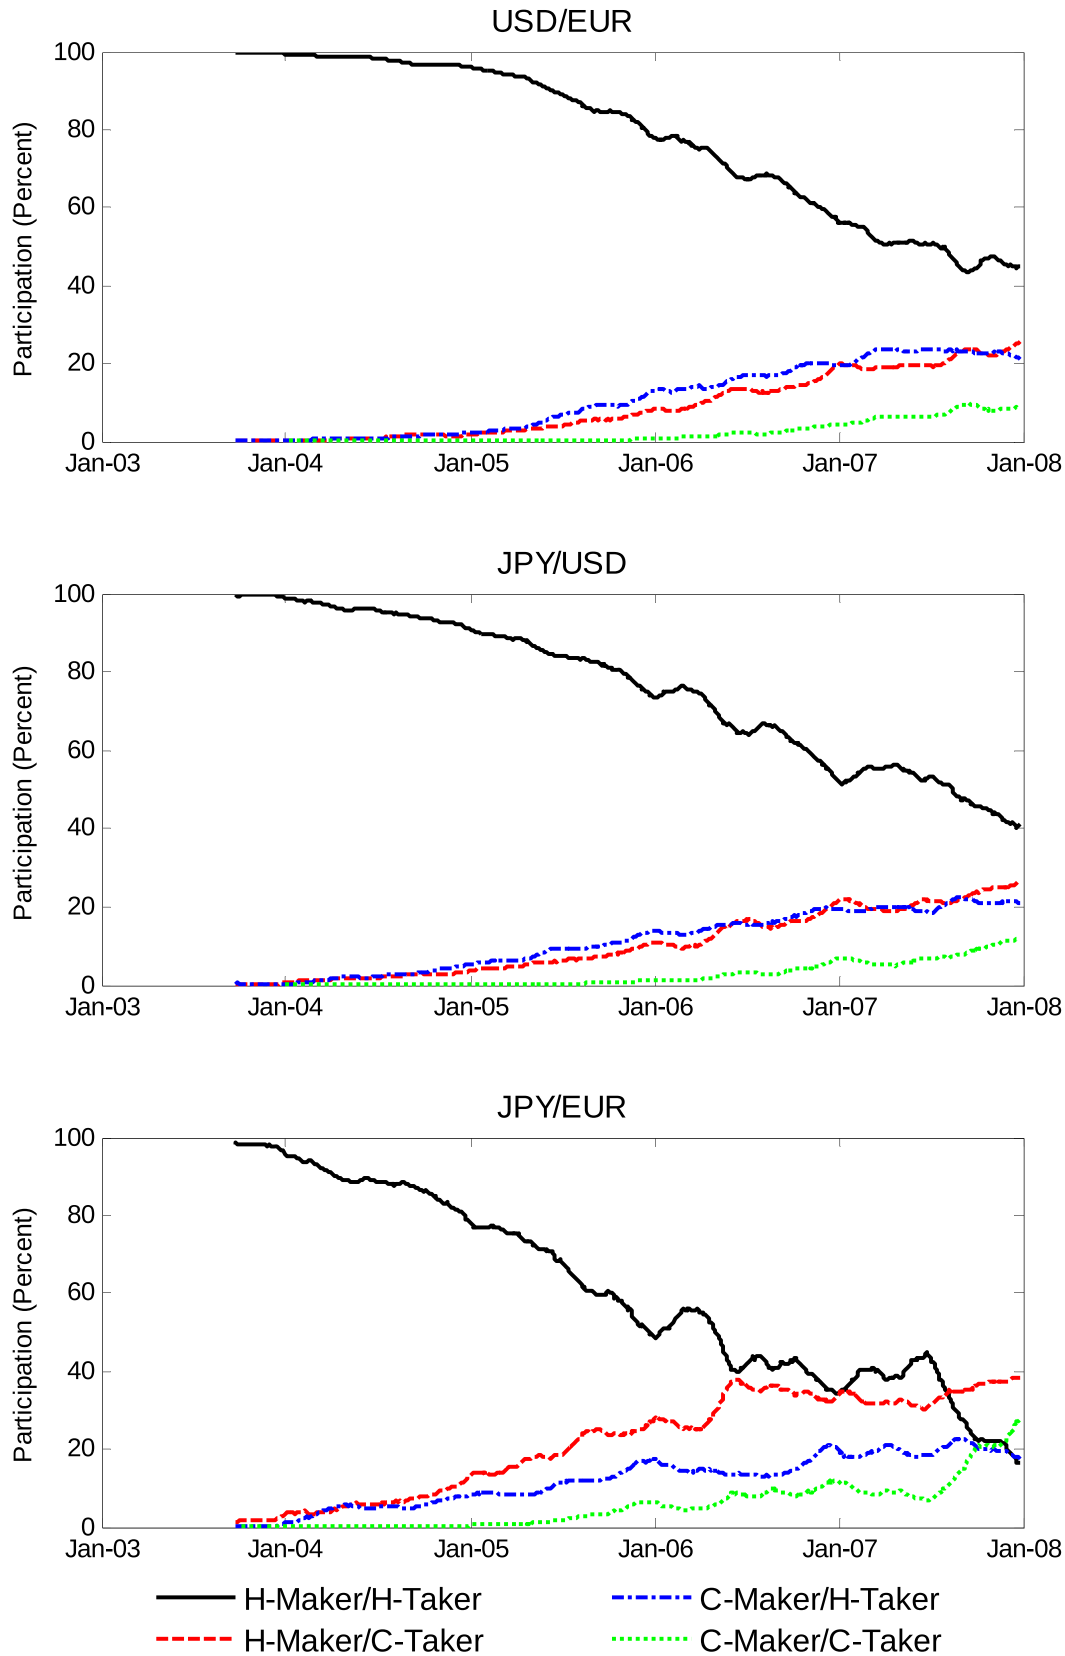

Figure 2: 50-day moving averages of participation rates broken down into four maker-taker pairs

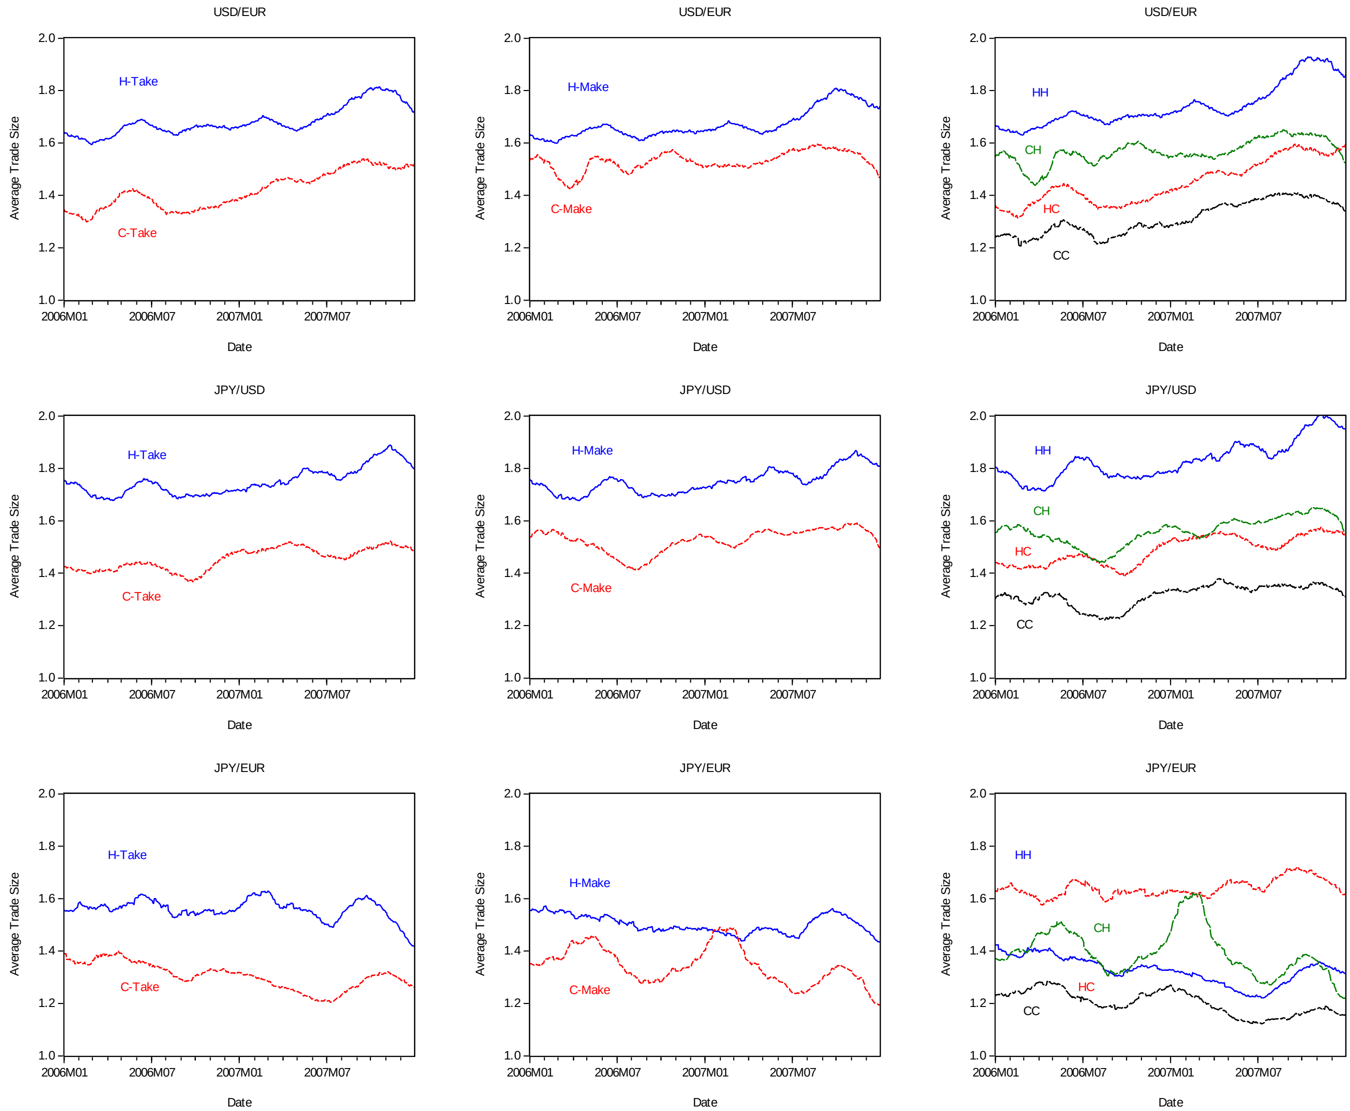

Figure 3: Average Trade Size

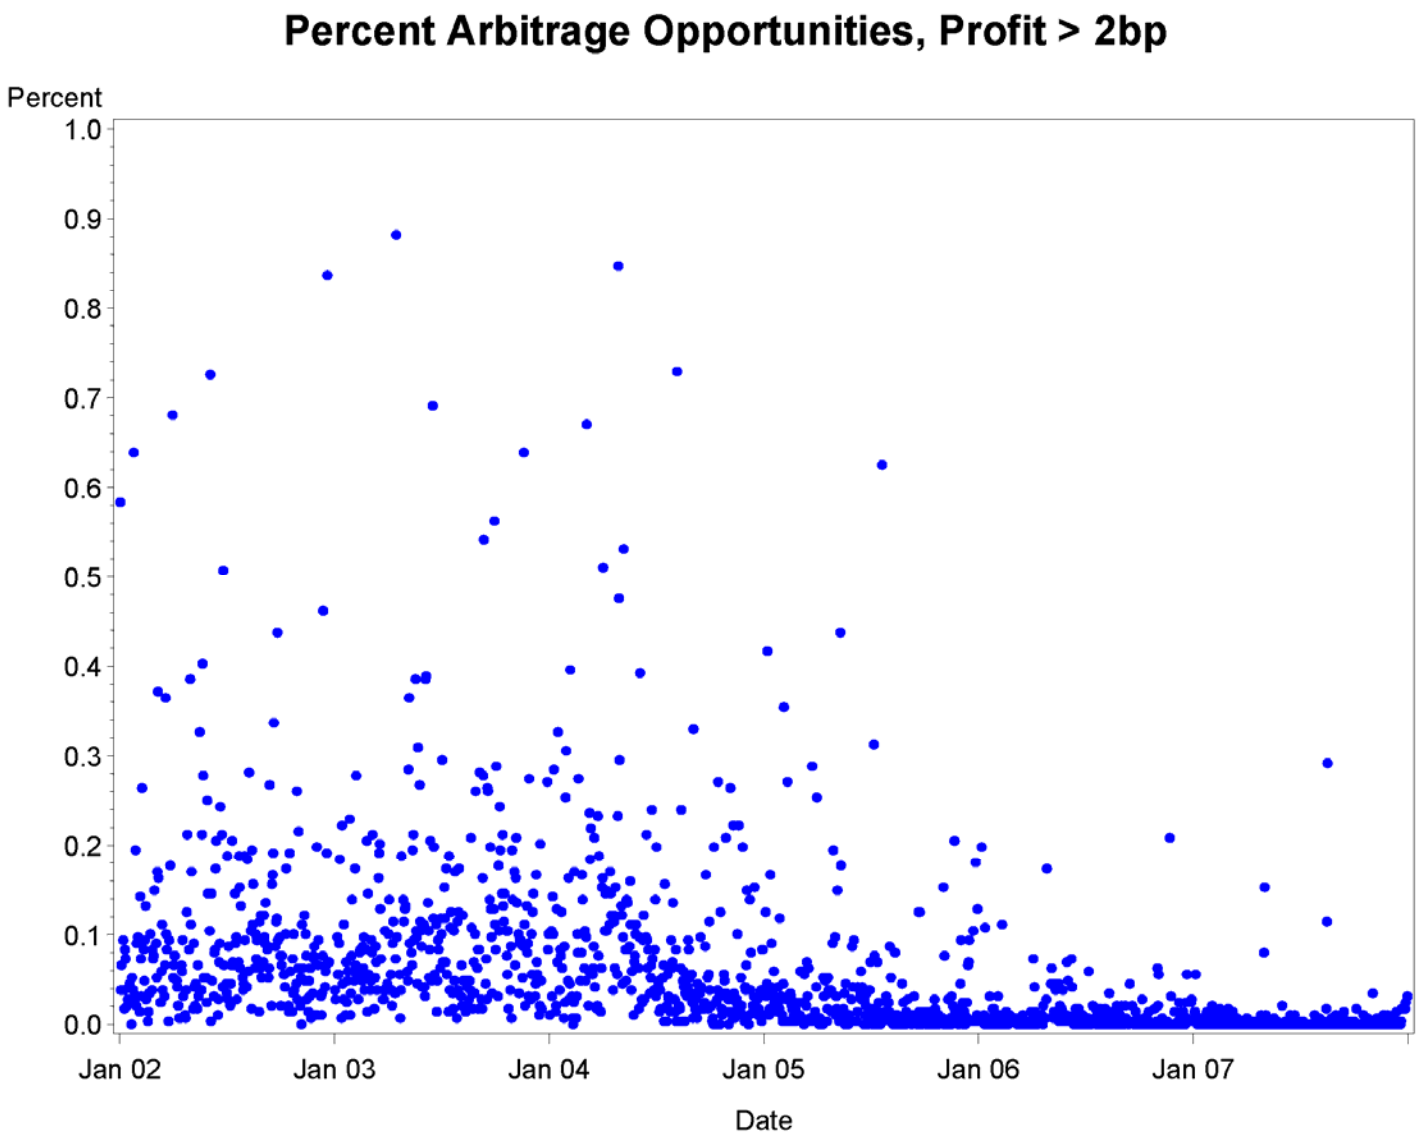

Figure 4: Triangular Arbitrage Opportunities

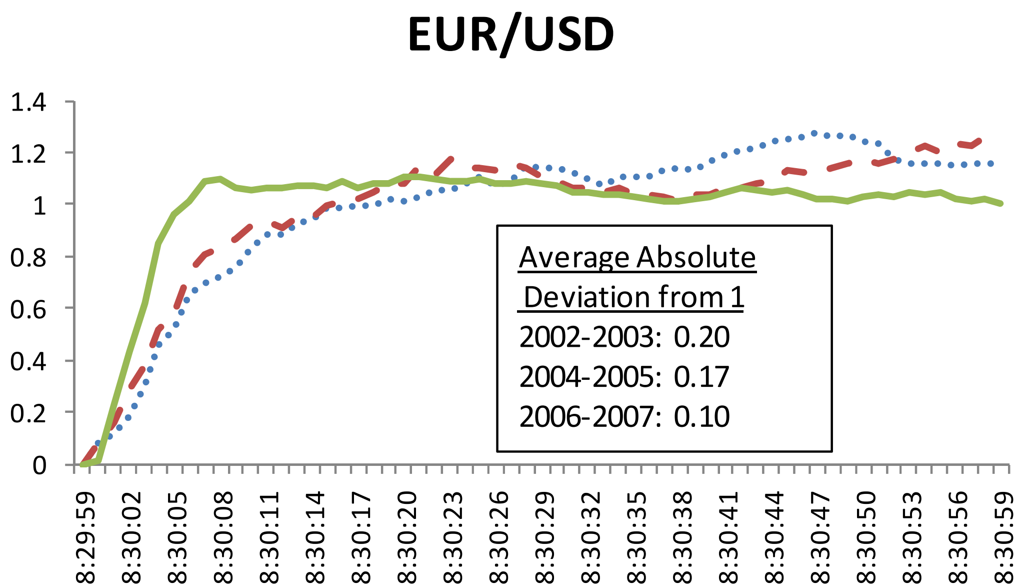

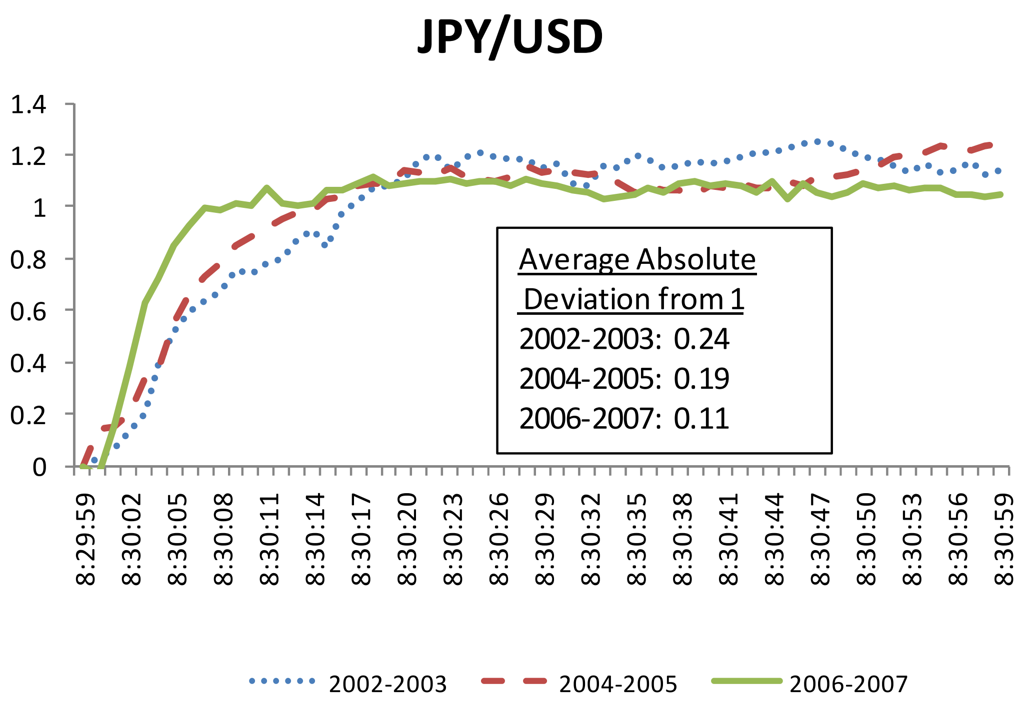

Figure 5: Response to non-farm payroll announcements

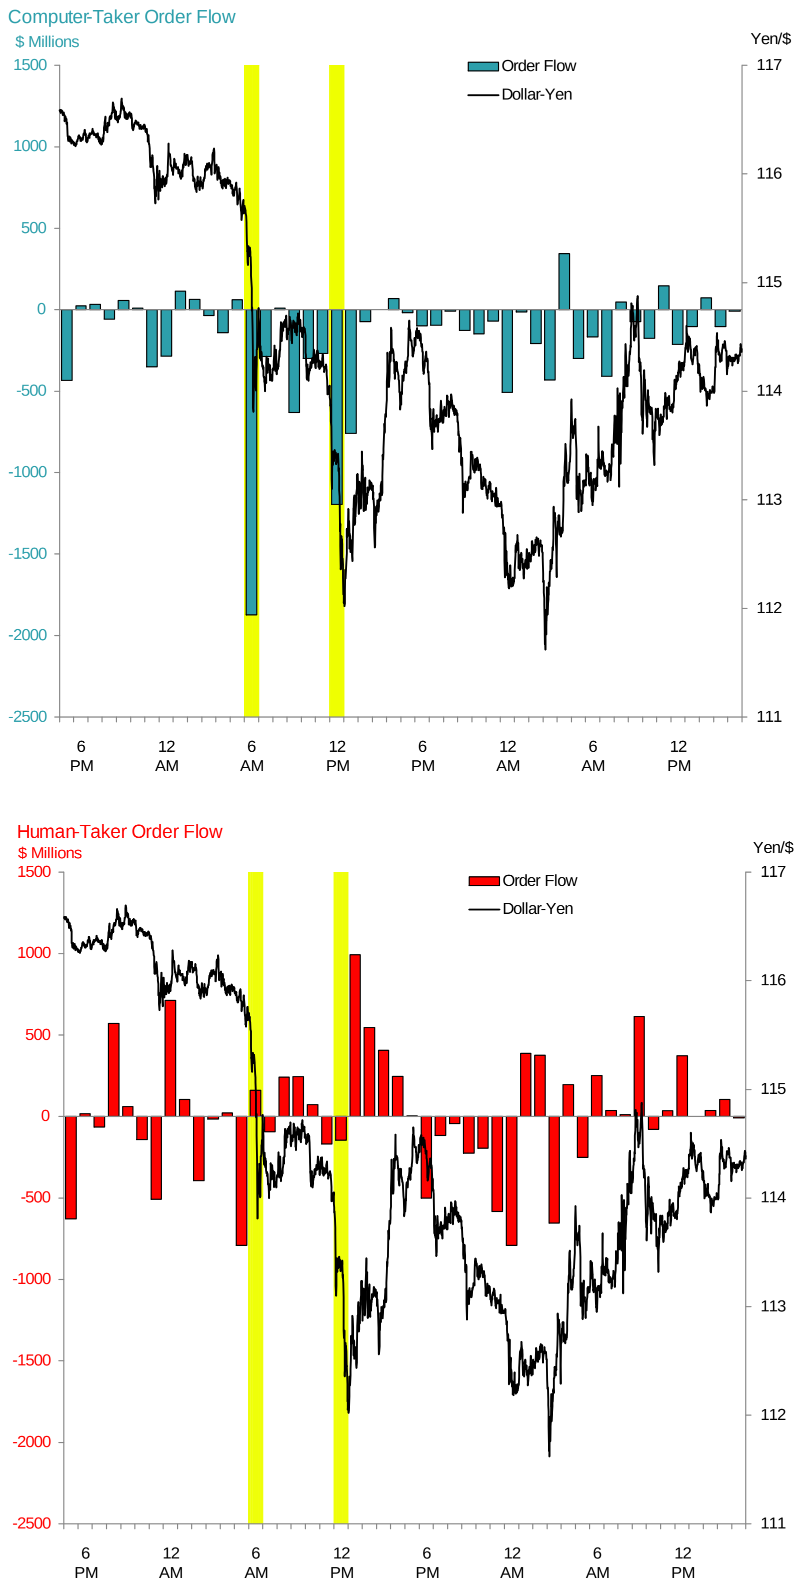

Figure 6: Dollar-Yen Market on August 16, 2007