The Microstructure of Foreign Exchange Dynamics

Martin D. D. Evans*
Department of Economics
Georegetown University
Washington DC 20057
Email evansm1@gunet.georgetown.edu
Phone (202) 687-1570
Fax (202) 687 6102

(First Draft: Comments Welcome)

December 4, 1998

Abstract

This paper studies the high frequency behavior of the interbank foreign exchange market with a newly created data set that provides the most comprehensive picture of activity across the market in existence. My analysis indicates that trade activity within the interbank market is distinct from the posting of indicative quotes. Trading and quote-making decisions are linked, but the links are complicated and poorly understood. I also document the existence of strong relationship between exchange rate movements and a measure of excess Dollar demand. A trading model is analyzed to show how the structure of the market could give rise to such a microstructure e®ect. Emprically, this e®ect appears important in the determination of exchange rates at high frequencies and over longer time spans relevant in international macroeconomics.

*Part of the research reported on in this paper was conducted while I was visiting the Financial Markets Group at the L.S.E. and The Bank of England. I am grateful for the hospitality I received at both institutions and acknowledge the financial support of the Houblon-Norman Fund at The Bank of England.

1 Introduction

This paper studies the high frequency behavior of the interbank foreign exchange (FX) market with a newly created dataset. This dataset has a number of unique features that allows us to examine the microstructure of the FX market in ways that were not previously possible. In particular, it contains time-stamped tic-by-tic data on both quotes and transactions for nine currencies throughout the interbank market from May 1 to August 31 1996. These data provide by far the most comprehensive picture of activity across the FX market in existence.

Most microstructure research on the FX market to date has utilized the quotes made by banks as shown on the screens of specialist information providers, such as Reuters, Telerate, Minex and Quotron.[1] These data can only tell us a limited amount about the microstructure of the market for several reasons. First, the quotes are only indicative. In other words, they do not represent the firm bid and ask prices at which a bank will enter into a transaction. Second, researchers have noted that as market activity intensifies traders become too busy to keep their indicative quotes up-to-date [Lyons (1996)]. Third, the quote series used by many studies is in fact a non random sample of the actually quotes posted by banks. Consequently, it is unclear whether the data in these studies can even be used to give accurate inferences about quote behavior.

In response to these shortcomings, a few studies have collected data on FX transactions. Lyons (1995, 1996) and Yao (1996) study the behavior of two individual dealers over 5 and 25 days respectively. Although the data in these studies are much more suitable for examining microstructure issues than the quote data, the degree to which they represent activity in the market as a whole is difficult to assess. Also, they are unsuitable for studying the role of heterogeneity in the workings of the market - an issue receiving increasing attention in the theoretical literature [see O'Hara (1994)]. As an alternative, Goodhart and Ito and Payne (1996) examine trading on a multilateral electronic trading system during a single day. Again, the degree to which these data are representative is unknown. As Goodhart, Ito and Payne acknowledge, trading in their sample may differ significantly from trading on the same system in the surrounding days and weeks. Moreover, transactions on the system are a tiny fraction of total trading activity so their trades may not be representative of market-wide activity even if the data came from a typical day.

This paper begins by describing the newly created data set in detail. I explain how the data was collected to make clear what information it can convey about the microstructure of the FX market. This discussion also explains why the quote data is not subject to the sampling problems that beset the quote series used in earlier studies.

[1] This literature has grown quite rapidly in recent years, thanks in part to the collection and dissemination of Reuter's quote data by the firm of Olsen and Associates. A partial list of papers includes; Baillie and Bollerslev (1991), Bollerslev and Domowitz (1993) and Bollerslev and Melvin (1994), Dacorogna, et al.(1993), Engle, Ito, and Lin, (1990), Goodhart and Giugale (1993), Guillaume, et al (1994a), and Ito, Engle, and Lin, (1992).

I next present a preliminary analysis of the data. The first noteworthy feature concerns the pronounced degree of heterogeneity across the currency spot markets. In terms of size and liquidity, the data indicate that the market for the DMark/Dollar is quite unlike other spot markets. As a consequence, the microstructure analysis of this particular market may not be very applicable to the spot markets for many other currencies. Nevertheless, my analysis focuses primarily on the DMark/Dollar in line with most existing research.

The seasonal patterns in transactions and quotes are examined next. Such patterns have been observed in previous studies, and are present here. I also study bivariate relationships between transactions and quotes. Since market-wide transactions data has not been available to earlier researchers, existing studies have tried to draw inferences about trading activity from the behavior of quotes.[2] My analysis provides some perspective on the accuracy of these inferences. I find that typical measures, such as the quantity of quotes and their volatility, only provide a moderately accurate indication of trade activity at high frequencies.

My analysis also considers the relationship between volatility, volume and spreads. The links between these variables in the FX market has been studied by numerous authors.[3] Here I find remarkably little evidence of links in either the transactions or the quote data. However, another relationship does stand out. I find that movements in transactions prices [i.e., the spot exchange rate] are strongly correlated with a market-wide measure of excess dollar demand. This relationship is present a very high frequencies and over long periods of days, weeks and possibly even months.

This finding is new to the microstructure literature on foreign exchange and is potentially important for two reasons. First, the result appears inconsistent with the standard efficient markets paradigm in which the institutional structure of the market is treated as irrelevant to the determination of equilibrium prices. Thus, the analysis of this new data allows us to make a premia facie case for relevance of market microstructure in the determination of exchange rates. Second, the statistical link between exchange rates and excess demand does not disappear when we consider the data over the longer time spans. Whatever its origins, this microstructure effect appears relevant to exchange rate dynamics at macroeconomic frequencies.

To examine the possible origins of the relationship between excess demand and prices changes, I develop a simple trading model. The model emphasis three features of the FX market; decentralized trading, the lack of transparency, and heterogeneous information amongst traders. It provides a relatively simple rationale for the correlation observed in the data.

The final section of the paper studies the dynamic properties of the transactions and quote

[2]For example, Bollerslev and Domowitz (1993) use a algorithm to construct artificial transaction prices from indicative quotes.

[3]See, for example, Bessembinder (1994), Bollerslev and Melvin (1994), Demos and Goodhart (1992), Glassman (1987), Hartman (1997), Guillaume et al (1994a), and Wei (1994).

data with a series of Vector Autoregressions (VARs). The aim here is two-fold. First, I want to document any characteristic features of the data that can serve as the basis for future theoretical research. Second, I wish to examine the extent to which high frequency dynamics of transactions data are reflected in the behavior of quotes, and vise-versa.

The results from the VAR analysis may be summarized as follows: There is ample evidence to suggest that the posting of quotes is a distinct activity rather than an adjunct of trading. Nevertheless, the data also indicate that quote and trading decisions are not make independently from one another. Changes in trading activity within the interbank market significantly affect quotes while innovations in quote activity affect transactions. These results suggest the presence of a complex interaction between transactions and quotes that needs to be explored theoretically. In short, quotes cannot be viewed a simply as noisy measures of transactions.

The paper is organized as follows. Section 2 describes the data. Section 3 presents the preliminary analysis of seasonality and the bivariate relationships between transactions and quotes. The trading model is presented and analyzed in Section 4. Section 5 examines dynamics with the aid of VARs. The paper ends with a brief summary.

2 Data Description

The analysis in this paper is based on a newly constructed micro data set detailing activity in the spot foreign exchange market from May 1 to August 31, 1996. The data set contains time-stamped tic-by-tic data on both indicative quotes and actual transactions for nine currencies. For the purpose of this study, I will focus primarily on the Dollar/DMark market. These data were collected via an electronic feed that was customized for the purpose by Reuters.[4] The feed contained all the information available to FX traders subscribing to the Reuters' systems. This is the sources of the quote data. The feed also contained transactions data that was unknown to market participants. Apart from a single power failure that interrupted data collection for one trading day, the data were collected continuously from May 1 to August 31, 1996.

2.1 Transactions Data

The transactions data come from the Reuters Dealing 2000-1 system. This is the most widely used electronic dealing system among foreign exchange dealers. According to Reuters, approximately 60% of transactions in the interbank market take place through the system. Unlike the electronic brokerage D2000-2 system examined by Goodhart, Ito and Payne (1996), D2000-1 is a bilateral

[4] This electronic feed was estimablished while I was visiting The Bank of England as the Houblon Norman Fellow. I am very greatful to both The Bank and Reuters for authorizing the feed and the help of numerous technical staff in setting it up.

trading system. Trades on this system take the form of electronic conversations. The conversation is initiated when a trader calls another trader on the system asking for a quote. Users of the system are expected to provide a fast quote with a tight spread which in turn are accepted or declined very quickly [approximately within three seconds]. Acceptance of a quote constitutes a trade. All electronic conversations, including trade confirmations, are printed out by the banks involved for record. In addition, Reuters keeps a record of all the conversations on the system for a limited time to settle disputes. This is the source of the transactions data. Every time an electronic conversation on D2000-1 results in a trade, the Reuters feed provided a time-stamped record of the transactions price, a bought or sold indicator, and a measure of cumulative trading volume. Reuters was unable to provide the identity of the trading partners for confidentially reasons.

Several features of the data are noteworthy. First, they provide information about transactions across the whole interbank market. It is clear from the daily pattern of trading intensity analyzed below that transactions are recorded between traders around the world. Transactions on the D2000-1 system provide the most comprehensive data source of market-wide trading activity [particularly in Europe and the US]. This contrasts with existing studies of individual traders by Lyons (1995, 1996) and Yao (1996), and the electronic brokerage D2000-2 system examined by Goodhart, Ito and Payne (1996) which captures a very small fraction of daily trading volume.

Second, these transactions data are not known to FX traders even though they trade on the D2000-1 system. While dealers have access to their own trading records, they are unable to receive a record of other transactions on the system. The transactions data therefore represents a history of market activity that can only be inferred indirectly by market participants.

Third, the data cover a reasonably long time span; 79 days after accounting for weekends, holidays and the feed interruption. This is important because it allows us to estimate "seasonal" patterns in market activity with much more precision than would be possible if the data only spanned a few days. Many existing studies of quote data have noted the importance of controlling for these "seasonals" when analyzing other features of the data [see, Baillie and Bollerslev (1991), Andersen and Bollerslev (1994), Dacorogna et al (1993) and Payne (1997)]. The span of the data also covers a period during which macroeconomic conditions around the world varied. As such the data has the potential to uncover the mechanism through which macroeconomic news impacts upon the market. Currently, this mechanism is very poorly understood. A fact that motivates this line of FX research.

2.2 Quote Data

The transactions data are complemented by the DFX series of indicative quotes reported on the Reuters information terminals. For each currency, the DFX series includes the time-stamped indicative bid and ask price with the identity of the bank posting the quote to Reuters. 526 different banks post quotes during the sample. For heavily traded currencies like the DMark/Dollar, it is not uncommon for several quotes to be reported each second.

The DFX series differs from the quote series that has been used in other studies, the FXFX series [see, for example, Baillie and Bollerslev (1991), Goodhart and Figliuoli (1991), and Guillaume et al (1994)]. Reuters transmits the FXFX and DFX data to its terminals through different networks. While the DFX data uses a high speed digital network that can keep up with the high frequency of quotes received from banks, the FXFX data uses an older system with a much lower throughput. To keep the FXFX series "up-to-date", Reuters only sends a sample of the quotes it receives down this network. As the frequency with which banks post quotes rises, Reuters has to send a smaller fraction of the quotes into the FXFX network to keep timely quotes arriving at terminals. Consequently, the FXFX data is a non-random sample of the quotes being posted to Reuters. This creates problems for the analysis of the FXFX data at very high frequencies. For this reason, the analysis below uses the DFX quote series in conjunction with the transactions data.

3 Preliminary Analysis

This section summarizes the main statistical features of the transactions and quote data. Since all the series at time-stamped to the second, the dataset is effectively a continuous time panel comprising the anonymous transactions data and the indicative quotes made by 526 banks. Analyzing this data is complicated by the fact that successive observations on each series occur at time-varying intervals. Moreover, since the timing of observations depends on market participants actions, it is unlikely to be independent of other variables. For example, the frequency of transactions may well by correlated with the variance of quotes measure over some past interval. Although models exist that can account for such dependencies, [see, for example, Engle (1996)], they require a good deal of structure. As such, they a not well-suited to the preliminary analysis undertaken here. Below I conduct my analysis over a fixed time interval of five minutes. The series studied will be statistics derived from the continuous data during this interval. Also, since my aim is to summarize the main statistical features of the market, I will not analyze the quotes made by individuals banks.

Table 1 shows the currencies covered by the dataset ranked by the number of transactions. Even at this highly aggregate level, several features are worth pointing out. First, there is a pronounced degree of heterogeneity across the individual spot markets. Whether measured by transactions or quotes, the DMark appears by far the most active market. By comparison, the markets for the Gilder and Krona appear very inactive. Of course these differences may not be a completely accurate reflection of market activity because not all bilateral transactions pass through the D2000-1 system. Nevertheless, the data in column (i) do suggest that the microstructure analysis of the DMark market is unlikely to be applicable across a large number of currencies.

Table 1: Data Description								
Currency/$	Trades	Quotes (DFX)	Ratio	Number of Banks				
				5	10	25	50	Total
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
German DMark	257,398	639,216	40.26%	28.62%	45.04%	70.32%	87.21%	382
Japanese Yen	152,238	163,074	93.36%	30.30%	47.70%	74.78%	87.73%	263
Swiss Franc	67,985	171,230	39.70%	52.38%	68.26%	91.43%	97.09%	194
Pound Sterling	52,318	137,656	38.01%	48.26%	69.78%	86.07%	94.17%	32
French Franc	20,553	75,219	27.32%	63.56%	80.43%	94.86%	99.43%	154
Italian Lira	8,466	165,144	5.12%	58.04%	79.40%	95.84%	99.25%	30
Begium Franc	5,256	36,143	14.54%	87.28%	95.61%	99.90%	99.99%	83
Dutch Guilder	3,646	129,450	2.82%	77.37%	91.68%	99.33%	99.99%	88
Danish Krona	1,488	72,037	2.07%	80.55%	97.53%	99.99%	99.99%	48

Notes: Columns (i) and (ii) report the total number of trades and quotes for each currency recorded from May 1 to August 31, 1996. Column (iii) reports the ratio of trades to quotes. Columns (iv) to (vii) shows the percentage of quotes originating from the 5 to 50 most active quoting banks for each currency. The total number of banks quoting the currency are reported under (viii).

The second striking feature of Table 1 concerns the ratio of transactions to quotes shown in column (iii). This ratio is remarkably similar amongst the most active European currencies; approximately 39%. The ratio for the Yen, by contrast, is over 93%. The most likely explanation for this difference relates to the international prominence of Reuters' information systems. In Asian markets, Reuters appears to face much stronger competition from other news services such as Minex than from alternative bilateral dealing systems. As a result, the Reuters system over-reports the ratio of Yen transactions to quotes during Asian trading. Amongst the less active currencies, the ratio of quotes to transactions is more varied, but it always appears to be well below the 40% level.

The right hand columns of the table present statistics on the distribution of quotes amongst banks. Columns (iv)-(vii) reports the fraction of quotes originating from the 5-50 banks that most actively quote the currency. These statistics indicate the degree to which quote activity is concentrated in a few banks. As the table shows, quote activity is heavily concentrated amongst a few banks particularly for the less actively traded currencies. In the case of the DMark, for example, 87% of the quotes originate from just 13% of banks that ever quote the currency [reported in column (viii)]. These heavy concentrations complicate inferences drawn from the quote data. They suggest that the information content of a quote may depend on both the quoted prices and the bank posting the information. As a result, it may not be appropriate to treat the quotes as a homogenous series, [i.e., as if they originated from a single representative bank], as has been the standard approach in the literature. The heavy concentration of quotes also calls attention to the fact that we have a very incomplete understanding of why banks post quotes. A fact that I shall return to below.

3.1 Seasonals

I begin my analysis be examining the recurrent patterns in trading and quote activity that take place over the trading day. Such patterns are often referred to as "seasonals" in the literature. Hereafter, I shall focus my discussion on the behavior of the DMark market.[5]

Several studies have documented systematic patterns in the intra-day volatility of indicative quotes [see, for example, Engle, Ito and Lin (1990), Baillie and Bollerslev (1991), Bollerslev and Domowitz (1993), Goodhart and Giugale (1993) and Payne (1997)]. These patterns are thought to be indicative of the intra-daily patterns in trading volume insofar as they appear related to the opening and closure of FX markets around the world. Figure 1 shows the intra-day patterns of volatility in ask quotes and transactions prices. Volatility in each series is calculated from the standard deviation of the second-by-second observations over a five minute period. The solid line plots the average volatility for the particular period over the 79 trading days in the sample. The lower and upper dashed lines show the 10% and 90% bounds of the volatility distribution calculated over the sample. These bounds indicate the degree to which volatility at a given time differs across trading days.

Panel I plots the volatility in ask quotes against British Summer Time (BST). On average, volatility rises over the trading day, peeking around 1600 hrs. BST. This pattern is similar to that found in earlier studies based on the FXFX quote data and are commonly attributed to market openings and closures around the world [Andersen and Bollerslev (1994), Dacorogna et al (1993), and Payne (1996)]. As support for this view, Panel II plots the seasonal pattern in quote intensity [measured by the number of quotes per five minute interval]. Here we see that quote activity is heavily concentrated between 800 hrs. and 1700 hrs. BST when the European markets are open.

How do these seasonal patterns in quote volatility and intensity compare with transactions prices? Panels III and IV plot the volatility of asking prices and transaction intensity [measured by the number of dollar purchases per five minute interval]. Here we see some evidence of a "U-shaped" pattern in volatility between 800 hrs. and 1700 hrs. BST. Although the change in average volatility is small compared to the error bands, this seasonal pattern is more pronounced than in the quote data. It is also worth noting that the average volatility of transactions prices is higher that the volatility of quotes, particularly around 800 hrs. BST. The "U-shaped" volatility pattern corresponds to the seasonal in trading intensity. Panel IV shows that on average intensity rises after 800 hrs. BST, peeks around 1000 hrs., and then declines towards mid-day before peeking again in the afternoon before 1600 hrs. BST.[6] Overall, these plots indicate that the seasonal relationship between quote intensity and volatility is reasonably representative of the relationship

[5] All the analysis reported below was also conducted with the Yen and the Pound. To conserve space, I shall only refer to these results when they stand in contrast with the DMark.

[6] In the case of the Yen, there is also a peek in activity around 400 hrs. BST.

between transaction intensity and volatility.

Figure 2 compares the seasonal patterns in quoted bid and ask prices with the bid and ask prices at which transactions take place. Panel I plots the mean spread between the quoted ask and bid DMark price for Dollars over five minute intervals averaged over the sample of 79 trading days. The most prominent feature of these data is the rise in the mean spread late in the trading day between 2000 and 2200 hrs. BST. This feature is surprising because it occurs well after the majority of quote and trade activity has died out [see Figure 1]. Moreover, it does not appear to be related to an "end-of-the-week" effect. The mean spread on Fridays, plotted with the dashed line, follows the same pattern.

Panel II plots the average volatility of the quoted spread [calculated as the standard deviation over five minutes] together with the 10% and 90% sample bounds. Spread volatility remains stable during periods where quote and trading activity is high; between 1000 and 2000 hrs. BST. Outside of this period, there is much more time series variation in average volatility and variability at a given time across different days. Because these patterns are heavily influence by the quotes of individual banks, they cannot be viewed as representing market-wide phenomena.

The DFX quotes are widely recognized as being "indicative" rather than firm, and so do not represent any commitment by the quoter to trade at the quoted prices. As such, it is unclear that the seasonal patterns in the spreads shown in Panels I and II can be interpreted as if they were the "firm" spreads faced by traders in the market.[7] Unfortunately, we cannot study the behavior of the firm spreads directly in this data set because we only observer the firm bid or ask price offered to a caller that results in a trade. To avoid such sample selection problems at this early stage, I consider another statistic derived from the transactions prices that provides further evidence of seasonality.

Panels III and IV plot the range of transactions prices averaged over the sample of 79 trading days and the bands of 10% and 90% [the range is calculated as the difference between the largest and smallest purchase (sale) price divided by the number of purchases (sales) during each five minute interval]. In both cases, there is evidence of a "U-shaped" pattern in the average range during the period of high trading activity; 800 to 1800 hrs. BST. Interestingly, the U-shaped pattern extends outside the period of heavy trading in the case of bid prices but not ask prices. At present, these patterns do not have any obvious interpretation.

[7]Nevertheless, several studies have examined their behavior in search of liquidity effects [see, for example, Glassman (1987), Bessembinder (1994) and Hartmann (1997)]. Other researchers, notably Lyons (1995) Goodhart, Ito and Payne (1996), have expressed reservations regarding the use of indicative spreads as a measure of the spread between firm quotes offered in the interbank market.

3.2 Quotes and Transactions

Most research to date has used indicative quote data to draw inferences about behavior in the interbank FX market. To examine the accuracy of these inferences, I shall analyze the relationships between the quote and transactions data at an hourly frequency with a series of scatter plots. These plots provide a simple non-parametric means to examine bivariate relationships in the data.

Figure 3 examines the links between quote intensity [measured as quotes per hour], volatility, trades, and trade volume. Panel I plots the relationship between quote intensity and the volatility of ask quotes calculated as the standard deviation of the last ask quote in each five minute period during the hour. As the plot shows, there is no evidence of a significant positive relationship between quote intensity and volatility measured at an hourly frequency. Similarly, Panel II shows that quote intensity is unrelated to the volatility of transaction prices [measured as the standard deviation of the last ask price in each five minute period during the hour]. These findings raise some questions as to the usefulness of quote intensity as a measure of market activity. Whatever motivates individual banks to change the frequency with which they post quotes, it does not appear to greatly affect the prices that are quoted [relative to other banks], or to reflect changes in the dispersion of transactions prices they observe during trading.

The lower panels of Figure 3 examine the links between quote intensity and two direct measures of trading activity. Panel III plots the relationship between quote and trade intensity. Although there variables are positively correlated, there is little evidence of a simple one-to-one the relationship. Rather the dispersion of trade intensity associated with a given number of quotes per hour appears larger at higher levels of quote activity. This cone-shaped scatter pattern is also evident in Panel IV which plots quote intensity against the dollar value of transactions during each hour. Clearly, observations on quote intensity do not provide a very accurate indicator of actual transactions activity.

The relationship between the volume of transactions and the volatility of transactions has been the subject of numerous studies investigating the process of price formation. For example, Ito, Lyons and Melvin (1997) examine how the introduction of trading in Tokyo over the lunch hour affects the volatility of FX quotes with the aim of discerning the importance of private information. This data set allows us to examine the impact of trading activity on other variables directly - we need not focus on the opening and closing of local markets.

Figure 4 presents scatter-plots examining the links between trade intensity, volatility, spreads and volume. Panels I and II show the relationship between trade intensity, the volatility of ask quotes, and the volatility ask transaction prices. These plots are quite similar to the volatility plots in Figure 3. Overall, they show no strong relationship between trade intensity and either volatility measure. While volatility does appear to rise with trade intensity when the latter is above 600 trades per hour, there are too few observations to draw any reliable inferences. Panel III compares trade intensity to the mean quote spread during each hour. Once again, it is very hard to see any relationship from the scatter-plot.[8] The relationship between trade intensity and volume is shown in panel IV. In contrast to the other plots, here there is quite a tight link between the variables. The R^2 statistic from a regression of volume on trade intensity is 0.87. The average size of a trade, measure by the regression slop coefficient is $3.6m. There is no regression evidence of a nonlinear relationship between trade intensity and volume - a finding consistent with the visual evidence in the scatter-plot. Consequently there is no indication in this data that the size of a typical transactions varies with trading intensity.

So far we have seen that quote intensity does not provide a very accurate indicator of transactions activity. Panels I and II of Figure 5 examine the accuracy of quoted prices. Panel I plots the volatility of ask transactions prices against the volatility of the ask quotes. Here we see that the volatility of quotes and transactions prices are only closely associated during periods of higher than average volatility. Otherwise, there is no clear relation. Thus, under most circumstances, changes in quote volatility are not indicative of changing volatility in transactions prices. Panel II compares hourly changes in quoted asking prices against changes in ask transactions prices. As the plot shows, this measure provides a much more accurate indicator of transactions prices than does quote volatility. Nevertheless, changes in the quoted prices are far from a completely accurate measure of changing transactions prices. Regressing the change in quoted prices on the change in transactions prices, we obtain a slope coefficient of 0.761 with a standard error of 0.012 and an R^2 statistic of 0.69. Thus changes in indicative quotes tend to understate the degree to which transactions price changed in the market. Panels III and IV compare trading volume against the quote spread and volatility. These plots are similar to those in Figure 4 based on trade intensity.

In summary, two features emerge from Figures 3 - 5. First, the behavior of both quoted and traded prices seems largely unrelated to quantity measures of quotes or transactions. Second, the quote data only provides a moderately accurate indication of the prices and quantities that arise from trade in the market.

These findings should not be viewed as too surprising. After all, the quotes are indicative and are widely recognized as being so both inside and outside the FX market. As such, it is unlikely that the decision to post a particular quote is tightly coordinated with a banks immediate trading activity. Moreover, even if some coordination is present, quotes can only reflect a very partial view of trading activity based on a bank's individual market experience. There is no reason to expect that quotes will accurately mirror transactions across the market at the high frequency studied here. In short, therefore, these results serve to underline the fact that trading and the posting of

[8] If the quotes were firm rather than indicative and so represented a commitment to trade, this plot would provide quite a challenge to theoretical models that link spreads to market liquidity. Here the absence of any relation most probably reflects the indicative nature of the quotes.

quotes are distinct activities.

3.3 Prices and Transactions Flows

One of the most interesting features of the transactions data is that is provides information about the aggregate volume and direction of transactions across the whole spot market. In particular, the data set contains information on whether the initiator of a transaction [i.e. the caller] bought or sold the Dollar. From these data we can construct high frequency measures of market-wide demand for currencies that provide some insight into market mechanisms that has not been possible hitherto. As noted in the introduction, existing studies of FX transactions have either been based on individual traders as is Lyons (1994, 1995) and Yao (1996), or electronic trading systems as in Goodhart and Ito and Payne (1995). In both cases the data used represents a tiny fraction of market-wide activity.

Figure 6 provides some preliminary analysis of these data. Panels I and II compare the daily behavior of the DMark/Dollar with a measure of excess Dollar demand[9],

$$D_t \equiv \frac{(\text{Number of Dollar Purchases}) - (\text{Number of Dollar Sales})}{\text{Number of Dollar Purchases + Sales}}.$$

Panel I plots the daily change in the ask transactions price against the excess demand measure. As the plot shows, there is a surprisingly strong and statistically significant positive correlation between these variables. On days when calls between traders resulted in many more dollar purchases than sales, the firm asking price for dollars rose. Panel II shows the cumulative effects of these demand imbalances on the log exchange rate over the sample period. The solid line shows that the daily changes in exchange rate cumulate into an sizable appreciation of the Dollar of almost 10%. The dashed line shows that the cumulative values of D_t follow a similar path.

The lower panels of Figure 6 examine the relationship between excess demand and price changes at the hourly frequency. Panel III plots the change in ask prices between the beginning of the day and the current hour against the D_t calculated in terms of the cumulated number of buys and sells to the current hour that day. As the plot shows, the relationship between these variables is more complex than in the daily data. Panel IV clarifies things by eliminating all the observations from Panel III that occur before 1200 hrs. BST. Here the remaining observations display the positive relation seen in the daily data.

Table 2 reports regression results that complement the visual evidence in Figure 6. The left

[9] Notice that this measure does not account for the traders wishing to buy or sell dollars who were unwilling to trade at the quoted price. Nor does it account for the different dollar value of individual trades. Moreover, the meaure is only derived from transactions executed over the D2000-1 system and so is only based on approximately 60% of the spot trades in the market. Despite these limitations, D_t is probably as good a measure of market-wide excess demand that we can construct.

Table 2: Excess Demand Regressions

$x_t = \alpha + \beta D_t + \sum_j \alpha_j I_{j;t} + \sum_j \beta_j I_{j;t} D_t + u_t$

x =	Cummulative Hourly Change				Hourly Change			
	(i)		(ii)		(iii)		(iv)	
β	0.0165	(9.360)	0.0008	(0.5034)	0.0012	(3.7448)	-0.0005	(1.0544)
β_8			0.0117	(2.4662)			0.0007	(0.5370)
β_{10}			0.0271	(4.1287)			0.0067	(3.2497)
β_{12}			0.0547	(6.3757)			0.0060	(3.3854)
β_{14}			0.0949	(9.3322)			0.0104	(5.6895)
β_{16}			0.1408	(11.1506)			0.0144	(6.9942)
β_{18}			0.1877	(14.1358)			0.0079	(4.7649)
β_{20}			0.2026	(23.2632)			0.0008	(1.0875)
R^2	0.0483		0.3846		0.0081		0.0728	
See	0.0049		0.0034		0.0014		0.0014	
d.f.	1725		1711		1725		1711	

Notes: In columns (i) and (ii) x_t is the cumulative hourly change over the day calculated as $P_t - P_t^0$; where P_t^0 is the first ask transactions price of the day, and P_t is last ask transactions price in the current hour. D_t is calculated in terms of the number of Dollar purchases and sales to the current hour during the day. $I_{j;t}$ is an indicator variable that takes the value of one if the observation falls between $(j-2) \times 100$ and $j \times 100$ hrs. BST. In columns (iii) and (iv) x_t is the hourly change in the ask transactions price and D_t is calculated in terms of the number of Dollar purchases and sales during the past hour. T-statistics are reported in parenthesis.

hand panel shows results from regressing the cumulative hourly change in asking price over the day, on cumulative hourly excess demand during the day. The positive and statistically significant slope coefficient reported under (i) confirms the visual evidence in Panel III of Figure 6. Under (ii), the table shows results from a regression that includes dummies for two hourly intervals in both the slope and intercept coefficients that run from 800 to 2000 hrs. BST. Here we see that the cumulative change in the spot rate is more strongly associated with the cumulative excess demand as the trading day progresses.

The right hand columns of the table report results from regressing the hourly change in spot prices on cumulative excess demand during the past hour. These regressions show us the degree to which variations in excess demand are associated with spot price changes at an hourly frequency. Under (iii) we see that there is a positive and statistically significant relationship - although it is weaker than that reported under (i). Column (iv) reports the results with seasonal dummies on both the intercept and slope coefficients. It is evident from these results that excess demand is most strongly associated with price between 1200 and 1600 BST, the period during which the market is most active [See Figure 1].

The results in Figure 6 and Table 2 are quite surprising. They suggest that typical traders systematically “under-adjust” their firm quotes to changing market conditions: To understand this interpretation, consider the perspective of a typical trader, Trader A, who receives a call from Trader B asking for a quote. Under normal circumstances, [i.e, in the absence of a large undesired inventory], we would expected Trader A to give quotes that equalize the ex ante probability of being hit at the bid or the ask. Although A knows the identity of B, he does not know B's reservation bid and ask prices which are based on her inventory, customer order flow and trading history. Instead, A must make an assessment of these reservation prices based on his own market experience.

Now suppose that customer orders around the market have generated an excess demand for dollars during the day. In other words, across the market more [firm] quotes have been hit at the ask than at the bid. No individual trader observes this directly. Instead their trading history contains a sample of this activity. If A's trading history contained a representative sample of market-wide activity under these circumstances, he would raise the ask price [relative to the bid price] to equalize the ex ante probability of being hit at either quote. Assuming A's actions are representative, their market-wide effects are clear. Although some callers will hit the higher asking price for idiosyncratic reasons, the number of calls across the market resulting in buys should fall relative to the number of sales. As a result, the cumulative imbalance between purchases and sales will fall. At the same time, the average sales price across the market will rise because some callers continue to hit the higher ask price. According to this logic, excess demand should not be positively correlated with the change in transactions prices because traders adjust their quotes in response to perceived market-wide pressures.

By contrast, the results above suggest that traders systematically under-adjust their quotes so that excess demand continues to build while quoted prices rise. This behavior appears hard to rationalize in centralized market with perfect transparency. As such, the results above make a prima facie case that the microstructure of the FX market, particularly its lack of transparency, is empirically relevant for understanding exchange rate behavior. Moreover, as Figure 6 clearly demonstrates, transaction imbalances appear to be associated with spot rate movements not just at very high frequencies but also over longer periods of days, weeks and possibly even months.

4 A Trading Model

To place further perspective on the statistical results presented above, this section analyzes a simple trading model. This model emphases the decentralized nature of FX trading, the lack of transparency, and the heterogeneity of trader's information. While I view these to be essential features of the FX market, the model below makes several simplifying assumptions and so does not account for all these features fully. In particular, it abstracts from the role played by inventory imbalances on traders' decisions. As such, the model cannot speak to the phenomena of “hot-potato" trading as analyzed by Lyons (1997). Nevertheless, it does provide a tractable analytical framework for studying the results above.

4.1 Traders

The market is populated by N traders who trade Dollars and DMarks bilaterally with one another. As in the FX market, bilateral conversations between traders are private information whether they result in a transaction or not. Thus at time t, each trader has a private history of quote and trade information as a result of the calls he received and made during the trading day until t. Traders also receive currency orders from customers outside the market [or possibly from other traders within the same bank]. I shall denote trader i's private information set at t containing both these information sources by V_t^i. In addition, traders have access to a common information set, Ω_t, that comprises current and past macroeconomic news, data and other publicly available information.

Traders actions depend upon their subjective assessment of the dollar's value. To formalize this idea, define P_t^* as the log DMark price of Dollars that would obtain in an economy with perfect information. Trader's decisions are based on their subjective probability distributions over P_t^*. In particular, I assume that traders actions depend upon their estimate of P_t^*, $E[P_t^*|V_t^i, \Omega_t]$. Differences in private information are assumed to be the only source of differences between traders' estimates of P_t^*. Thus if traders i and j have identical private information [i.e., $V_t^i = V_t^j$], $E[P_t^*|\Omega_t, V_t^i] = E[P_t^*|\Omega_t, V_t^j]$. This restriction rules out the possibility that two traders would value the Dollar differently for idiosyncratic reasons unrelated to the information they posses.

In a general equilibrium model, the evolution of V_t^i will be determined by news of customer orders from outside the market and the history of calls and trades that are determined endogenously. To avoid the complexity of determining V_t^i endogenously, I assume that traders' expectations take the form:

$$E[P_t^*|V_t^i, \Omega_t] = E[P_t^*|\Omega_t] + E[\varepsilon_t|V_t^i], \tag{1}$$

where $\varepsilon_t \equiv P_t^* - E[P_t^*|\Omega_t]$ is the error in estimating P_t^* based solely on public information. $E[\varepsilon_t|V_t^i]$ is the private component of trader i's valuation of the Dollar. This component evolves according to

$$E[\varepsilon_t|V_t^i] = \rho E[\varepsilon_{t-1}|V_{t-1}^i] + v_t^i \qquad 1 > \rho > 0 \tag{2}$$

where v_t^i is an innovation that represents the revision in trader i's valuation due to private news. Such news can be correlated across traders. For example if a large number of customers simultaneously place orders to buy dollars through numerous traders, the $v_t^{i\prime}$s may be positively correlated. To allow for these market-wide correlations, I assume that v_t^i can be decomposed into two independent components:

$$v_t^i = v_t + \eta_t^i; \quad (3)$$

where v_t and η_t^i are i.i.d. normally distributed, mean zero, random variables that represent market-wide and idiosyncratic news respectively. Because customer orders are private information, individual traders cannot discern the source of innovations. Instead, they form inferences based on the relative variances of v_t and η_t^i, denoted by σ_v^2 and σ_η^2. In particular, trader i's estimate of the market-wide component is $E[v_t|v_t^i] = \theta v_t^i$ where $\theta = \sigma_v^2/(\sigma_\eta^2 + \sigma_v^2)$.

This specification of traders expectations is highly structured in order to facilitate the analysis below. Nevertheless, it captures three important features. First, traders cannot tell whether the news they receive from customer orders or inter-dealer calls about the value of the Dollar is being simultaneously received elsewhere. This is clearly a feature of the FX market and has important consequences as we shall see. Second, private news has some persistent effect on the trader's valuation. Without some persistence, disagreement amongst traders about the valuation of the dollar would disappear very quickly after customer orders stopped arriving at the market. This appears inconsistent with the heterogenous views concerning the value of the dollar expressed by traders at the end of the trading day. It is also hard to square with the heterogeneity of spot rate expectations observed by Ito (1990) in micro survey data. Finally, the structure of expectations implies that most disagreement about the value of the dollar is short-term. Assuming that traders do not have advanced warning of public news announcements,[10] we can iterate (1) and (2) forward h periods, and take conditional expectations to obtain,

$$E[P^*_{t+h}|v_t^i; \Omega_t] = E[P^*_{t+h}|\Omega_t] + \phi^h E[e_t|v_t^i].$$

As the horizon h increases, individual traders' valuations are less influenced by their private information (because $\phi < 1$) so there is less disagreement about the long-term valuation of the currency. This formulation roughly accords with the degree of consensus amongst economists about the factors affecting the behavior of the Dollar at different horizons.

4.2 Quotes and Trade

Conversations between traders are initiated by an exogenous matching mechanism. Each trader receives and makes a call every period with different counter-parties. In reality, traders can decide who to call over the D2000-1 system but they can receive quote requests from any caller. The exogenous matching mechanism abstracts from the trader's decision of who to call.

Consider a call from trader i to trader j. Trader j makes quotes of $Q_t^{B,j}$ and $Q_t^{A,j}$ corresponding

[10] This rules out any correlations between $E[P^*_{t+h}|\Omega_{t+1}] - E[P^*_{t+h}|\Omega_t]$ and v^i_{t-s}, for all s, $h \geq 0$.

to the bid and ask prices at which he will sell and buy Dollars. I assume that all quotes and transactions are for a fixed quantity. By convention, $Q_t^{B,j}$ and $Q_t^{A,j}$ are take-it-or-leave-it quotes. Trader i can decide to either "hit" the ask or bid price, thereby agreeing to buy of sell Dollars, or they can decline to trade. This decision is determined by the relationship of the quotes to trader i's reservation price:

$$R_t^i = E[P_t^* | V_t^i; \Omega_t] + u_t^i. \quad (4)$$

The reservation price varies from the private valuation of the Dollar by a component u_t^i that reflects a liquidity effect. This effect originates from customers outside the market who need currency as part of their normal international business activities. Some liquidity effects are common to all traders, others are purely idiosyncratic. That is,

$$u_t^i = u_t + \omega_t^i$$

where u_t and ω_t^i are the common and idiosyncratic components. Both components are assumed to be i.i.d. normal random variables with zero means and variances σ_u^2 and σ_ω^2. As such, they only have a transitory effect on reservation prices.

The decision to trade is determined as follows: If $Q_t^{A,j} \leq R_t^i$ then trader i hits the ask price and buys dollars from trader j for price $Q_t^{A,j}$. [Note that this transaction is recorded as a buy in the data set.] Similarly, if $Q_t^{B,j} \geq R_t^i$ then trader i hits the bid price and sells dollars to j in exchange for $Q_t^{B,j}$ per dollar. If neither condition is met, the conversation ends without a transaction. Notice that the purchase of dollars is more likely, ceteris paribus, when u_t^i is positive so u_t^i plays the role of a liquidity demand shock.

When trader j receives a request for quotes, he faces a complex decision. Because the caller's reservation price is private information, the choice of quotes must be made relative to an estimate of R_t^i. Given this estimate, trader j will, in general, choose quotes to maximize expected trading profits during the day subject to some inventory constraints. This means that the quotes made by each trader will depend on their private information and their trading history through their current inventory. Here I focus solely on the effect of private information by assuming that current inventory does not affect quotes. This assumption appears reasonable during the middle of the trading day where the overnight inventory constraints placed on typical traders are far from binding.[11] In particular, I assume that trader j sets quotes so that there is an equal ex ante probability that trade i will buy or sell. Following this strategy, a trader should not accumulate or loose inventory on average.

[11] This assumption is consistent with the analysis of a single trader by Yao (1997). He does not find any evidence that quotes are adjusted in response to inventory imbalances. Instead, these imbalances are unwound by initiating transactions with other traders.

To determine how quotes are set under these circumstances, we need to equate the ex ante probabilities of a purchase and sale from the perspective of trader j: These probabilities are

$$\Pr\left(\text{trader } i \text{ buys dollars}\right) \equiv \Pr\left(Q_t^{A,j} \le R_t^i | V_t^j, \Omega_t\right)$$
$$\text{and} \qquad \Pr\left(\text{trader } i \text{ sells dollars}\right) \equiv \Pr\left(Q_t^{B,j} \ge R_t^i | V_t^j, \Omega_t\right):$$

Trader j calculates the distribution of the reservation price, R_t^i, based on his own private news and liquidity shocks $\{v_{t-h}^j, u_{t-h}^j\}_{h \ge 0}$. Under the assumptions above, the conditional distribution of R_t^i given trader j's information is normal with mean and variance,

$$E[R_t^i | V_t^j, \Omega_t] = E[P_t^* | \Omega_t] + \psi u_t^j + \lambda E[e_t | V_t^j];$$

$$\sigma^2 = \sigma_u^2(1 - \psi^2) + \tfrac{1-\lambda^2}{1-\phi^2}(\sigma_v^2 + \sigma_\eta^2);$$

where $\psi \equiv \sigma_u^2/(\sigma_u^2 + \sigma_\omega^2)$: Using this result, we can write the ex ante trade probabilities as

$$\Pr\left(Q_t^{B,j} \ge R_t^i | V_t^j, \Omega_t\right) = \Phi\left(\frac{Q_t^{B,j} - E[R_t^i | V_t^j, \Omega_t]}{\sigma}\right)$$
$$\Pr\left(Q_t^{A,j} \le R_t^i | V_t^j, \Omega_t\right) = 1 - \Phi\left(\frac{Q_t^{A,j} - E[R_t^i | V_t^j, \Omega_t]}{\sigma}\right) = \Phi\left(\frac{E[R_t^i | V_t^j, \Omega_t] - Q_t^{A,j}}{\sigma}\right) \qquad (5)$$

where $\Phi(.)$ denotes the cumulative standard normal. Clearly these probabilities will be equal when quotes are set as

$$Q_t^{A,j} = E[R_t^i | V_t^j, \Omega_t] + s/2$$
$$Q_t^{B,j} = E[R_t^i | V_t^j, \Omega_t] - s/2 \qquad (6)$$

for some spread s.

In choosing the value of the spread, traders face a trade-off. Since the ex ante probability of trading is $2\Phi(s/2\omega)$; larger values of s reduce the likelihood or being "hit" at either quote. At the same time, higher values of s raise the expected steady state profit from trading. For example, consider trader j's expected profit from selling Dollars to a caller at t for $Q_t^{A,j}$ and buying them back from another caller at $t+1$ for $Q_{t+1}^{B,j}$: In the steady state [where $E[e_t | V_t^j] = 0$]; the expected profit from the trades is $E[Q_t^{A,j} - Q_{t+1}^{B,j} | V_t^i, \Omega_t] = s$. Since the probability of being hit at the ask and bid in successive periods is $\Phi(s/2\sigma)^2$; the expected steady state trading profit is $\pi(s) = s\Phi(s/2\sigma)^2$:

I assume that traders choose s to maximize $\pi(s)$. This gives a value for s of $\sqrt{2}\sigma$.[12]

We can now calculate the probability of a trade taking place. Recall that trader i will buy dollars when $Q_t^{A,j} \leq R_t^i$ and sell them when $Q_t^{B,j} \geq R_t^i$. Using the results above, we can rewrite these conditions as

$$\begin{aligned} \text{trader } i \text{ buys if} \quad & z_t \leq -s/2 \\ \text{trader } i \text{ sells if} \quad & z_t \geq s/2 \end{aligned}$$

where

$$z_t \equiv \psi u_t^j - u_t^i + \lambda E[e_t|V_t^j] - E[e_t|V_t^i].$$

z_t is a normally distributed random variable with mean $(\lambda - 1)\mu_t$ and variance σ^2 where $\mu_t \equiv \sum_{h=0}^{\infty} \phi^h v_{t-h}$ is the average value of $E[e_t|V_t^i]$ across traders. Hence the probability of trade is

$$\begin{aligned} \Pr\left(\text{trader } i \text{ buys}\right) &= \Phi\left(\frac{-s/2 + (1-\lambda)\mu_t}{\sigma}\right) \\ \Pr\left(\text{trader } i \text{ sells}\right) &= 1 - \Phi\left(\frac{s/2 + (1-\lambda)\mu_t}{\sigma}\right) = \Phi\left(\frac{-s/2 - (1-\lambda)\mu_t}{\sigma}\right) \end{aligned} \tag{7}$$

Notice that ex post probability of a purchase or sale are not equal except in the special cases where $\mu_t = 0$ or $\lambda = 1$ even though trader j's ex ante probabilities of being hit at either quote are equal. This difference between the ex ante and ex post probabilities is a reflection of asymmetry between the traders. Quotes are made without any knowledge of the caller's reservation price and so are based solely on trader j's information. In particular, j must base his quotes on an estimate of the market-wide component of $E[e_t|V_t^j]$, which is equal to μ_t. So long as there is a perceived idiosyncratic element in $E[e_t|V_t^j]$, λ will be less than one and j's estimate will be less than completely precise. Trader i, by contrast, has no need to estimate μ_t when making a trading decision because he can request quotes before deciding whether to trade.

During a conversation, information is exchanged between the traders in several ways. First, the quotes made by trader j reveal to trader i something about the j's liquidity shock u_t^j and $E[e_t|V_t^j]$. Given the form of the (6), there is insufficient information in these quotes to completely unravel u_t^j and $E[e_t|V_t^j]$. Nevertheless, a caller would find the quotes useful in updating his priors about μ_t and u_t. Thus the model provides an information collection rational for initiating a call. Second, trader i's decision to trade reveals information to the quoter. In particular, trading decisions provide j will some information on the value of z_t. Since $u_t^i + E[e_t|V_t^i] = \psi u_t^j + \lambda E[e_t|V_t^j] - z_t$, this, in turn,

[12] Outside the steady state, the value of s that maximizes expected profits depends on $E[e_t|V_t^j]$. In this case, the quotes would reveal all trader $j's$ private information to the caller.

allows j to update his priors about μ_t and u_t:

4.3 Aggregate Activity

We are now in a position to analyze trading activity across the market. Consider first the behavior excess dollar demand: During the middle of the trading day, where the number of traders is large, excess dollar demand, D_t; can be approximated by

$$D(\mu_t) = \frac{\Phi\left(\frac{-s/2+(1-\lambda)\mu_t}{\sigma}\right) - \Phi\left(\frac{-s/2-(1-\lambda)\mu_t}{\sigma}\right)}{\Phi\left(\frac{-s/2+(1-\lambda)\mu_t}{\sigma}\right) + \Phi\left(\frac{-s/2-(1-\lambda)\mu_t}{\sigma}\right)}; \tag{8}$$

which is increasing in μ_t, the market-wide average of $E[e_t|V_t^j]$: Equation (8) implies that the arrival of private news that leads traders across the market to raise [lower] their valuation of the dollar will generate greater [less] excess demand in dollar transactions.

Notice that movements in $D(\mu_t)$ only arise because λ is less than unity, a reflection of the fact that traders do not know the degree to which their private news is correlated with others. In this respect, the model points to an important implication of the lack of transparency in the FX market. To further emphasis this point, note that the arrival of public news during period t, i.e., $E[P^*_{t+h}|\Omega_t] - E[P^*_{t+h}|\Omega_{t-1}]$; has no effect on either the number or type of transactions, because, by definition, it is transparent to all market participants.

What are the model's implications for the behavior of transactions prices? To answer this question, consider the market-wide average of the price paid for Dollars during period t: From equation (6) the average asking price at t is

$$Q_t^A = E[P_t^*|\Omega_t] + \psi u_t + \lambda\mu_t + s/2 \tag{9}$$

where u_t is the common component of the liquidity shock affecting reservation prices. Individual quotes differ from Q_t^A by a random amount that is specific to the trader making the quote. These quotes are observed in the data only if the caller decides to buy dollars. Transaction prices are therefore a censored series of the ask quotes with the degree of censoring determined by the probability that a purchase results from a conversation. From equation (7) we see that this probability only varies with μ_t and not the asking price because the decision to buy depends on the difference between the ask and reservation price. This means that the sampling of quotes in the transactions data is independent of the transactions prices. We may therefore represent the observed purchase price as

$$P_t^A = Q_t^A + \xi_t; \tag{10}$$

where ξ_t represents the combined effects of sampling and the idiosyncratic features of the individual ask quote that was "hit". Importantly, ξ_t is independent of Q^A_t for the reasons cited above.

We can now examine the observed change in ask prices by combining (9) and (10). Taking first differences and simplifying, we obtain

$$\Delta P^A_{t+1} = \left\{ E[P^*_{t+1}|\Omega_{t+1}] - E[P^*_t|\Omega_t] \right\} + \lambda(\phi - 1)\mu_t + \lambda v_{t+1} + n_{t+1} \tag{11}$$

where $n_{t+1} \equiv \Delta\xi_{t+1} + \psi\Delta u_{t+1}.$

The first term in (11) identifies the impact of public news on transactions prices. This term is identically equal to the expected change in P^*_t and the revision in the expected value of P^*_{t+1} given the arrival of public news,

$$E[\Delta P^*_{t+1}|\Omega_t] + E[P^*_{t+1}|\Omega_{t+1}] - E[P^*_{t+1}|\Omega_t].$$

Since intra-day interest rates faced by traders are zero, $E[\Delta P^*_{t+1}|\Omega_t]$ identifies the expected excess return on Dollars over DMarks during the next period based on public information. This return is equal to the dollar risk premium and will generally be very small. As a result, the first term in (11) is well approximated by $\varepsilon_{t+1} \equiv E[P^*_{t+1}|\Omega_{t+1}] - E[P^*_{t+1}|\Omega_t].$ We can therefore rewrite (11) as

$$\Delta P^A_{t+1} = \lambda(\phi - 1)\mu_t + \lambda v_{t+1} + \varepsilon_{t+1} + n_{t+1}. \tag{12}$$

According to (12), price changes may be predictable for two reasons. First, since $\phi < 1$, ΔP^A_{t+1} will co-vary negatively with μ_t provided $\lambda > 0$. This means that price changes could in principle be predicted with μ_t. However, if we are considering short time periods, such as a hour, the appropriate value of ϕ is probably very close to one. As a result, unless μ_t has a very large variance, variations in ΔP^A_{t+1} are likely to be almost completely dominated by the last three terms. The second source of predictability comes from n_{t+1} because, under the assumptions of the model, it follows an MA(1) process. Once again, this is unlikely to generate significant forecastability in ΔP^A_{t+1} unless the variance of $\lambda v_{t+1} + \varepsilon_{t+1}$ is small. In sum therefore, (12) does not suggest that ΔP^A_{t+1} should be forecastable. A fact that appears consistent with the data.

Recall from Table 2 that the change in ask prices appear positively correlated with excess demand, particular towards the end of the trading day. We can now interpret this result with (8) and (12). As the equations show, the arrival of private news across the market, v_{t+1}, raises both transactions prices, P^A_{t+1}, and excess demand, D_{t+1}. By contrast, the arrival of public news, ε_{t+1}, only affects transactions prices. Since the announcement of public information is concentrated during the first half of the trading day, any positive correlation between P^A_{t+1} and D_{t+1} induced by v_{t+1} shocks during this period will likely be masked. Later on, public news announcements are much less frequent, so the positive correlation between P^A_{t+1} and D_{t+1} is revealed.

In summary, this simple trading model provides a rationale for the results in Table 2. In particular, we have seen that the "under-adjustment" of quotes in response to excess demand can arise quite naturally in a market that lacks transparency when traders have heterogeneous information. Here the lack of transparency stopped traders from observing the extent to which the arrival of private information was correlated across the market. This creates an informational asymmetry between traders during bilateral conversations that is the source of the correlation between transactions prices and excess demand.

5 Dynamic Analysis

In this final section, I return to the analysis of the data. I shall focus on the dynamic relationships between transaction prices, quantities and quotes. The aim of this analysis is two-fold. First, I want to document any characteristic features of the data that can serve as the basis for future theoretical research. Second, I wish to examine the extent to which the high frequency dynamics of transactions data are reflected in the behavior of quotes. As note earlier, there are now a large number of studies documenting the high frequency behavior of the quotes. However, we do not yet know the extent to which this behavior reflects transactions activity within the market rather than the decisions by banks to make quotes.

My analysis is conducted using VARs. In particular, since my focus is on the high frequency dynamics, the analysis is conducted on the data sampled every five minutes during the trading day. In view of the seasonal patterns in Figure 1, the trading day is defined by the period between 700 and 1700 hrs. BST. The VARs include a set of dummy variables to allow for the effect on non-continuous observations, such as between trading days.

5.1 Trade and Quote Quantities

The first VAR examines the relationship between the number of trades, quotes and the volatility of transactions prices. Recall from Section 3 that at an hourly frequency the contemporaneous relationship between trade and quote intensity appears quite loose. We also saw that the link between transactions intensity and the volatility of transactions prices was week. The VAR permits us to investigate these relationships in greater detail.

Table 3 shows summary estimation results for a 6'th-order tri-variate VAR using the variables, $\{NT_t; NQ_t/NT_t; \sigma(T)_t\}$ where, NT_t and NQ_t are the number of trades and quotes and $\sigma(T)_t$ is the standard deviation of ask transaction prices during each five minute period. Panel I reports p-values for a series of Granger Causality tests. There is strong evidence of causality running in all directions except in the case of the volatility equation. Here we cannot reject the null of no Granger Causality running from trade intensity, NT_t or from relative quote intensity, NQ_t/NT_t to the volatility of transactions prices at the 5% level.

Table 3: VAR results for Trade, Quote and Trade Volatility							
I Granger Tests							
VAR Equation		NT_t		NQ_t/NT_t		$\sigma(T)_t$	
	NT_t	<0.001		<0.001		0.299	
	NQ_t/NT_t	<0.001		<0.001		0.124	
	$\sigma(T)_t$	<0.001		<0.001		<0.001	
II Variance Decompostions							
horizon	source	%	(Std. Err.)	%	(Std. Err.)	%	(Std. Err.)
1 hour	NT_t	88.100	(6.036)	13.009	(6.742)	40.707	(22.040)
	NQ_t/NT_t	10.498	(6.185)	80.397	(7.467)	42.900	(22.205)
	$\sigma(T)_t$	1.402	(1.537)	6.595	(2.992)	16.392	(9.033)
3 hours	NT_t	59.197	(22.200)	0.004	(9.585)	10.313	(31.036)
	NQ_t/NT_t	40.770	(22.701)	95.679	(10.553)	82.967	(31.381)
	$\sigma(T)_t$	0.033	(1.409)	4.317	(2.917)	6.719	(3.387)
III: Diagnostics							
Q_3		0.989		0.997		0.997	
Q_6		0.673		0.864		0.990	
Eigenvalues		0.887		0.857		0.635	

Notes: NT_t and NQ_t are respectively the number of trades and quotes and $\sigma(T)_t$ is the standard deviation of ask transaction prices, each calculated over a five minute period. The variables are ordered in the VAR as $\{NT_t; NQ_t/NT_t; \sigma(T)_t\}$. The VAR contains 6 lags of each variable as well as a set of dummy variables to allow for discontinuities in the data. The dummies take the value of one if observations x_t and x_{t-1} are not based on data from two consecutive five minute intervals, and zero otherwise. Observations are only taken between 7.00 and 17.00 hrs. BST. Q_i denotes the p-values for Box-Pierce Q statistics for i'th. order residual serial correlation in each VAR equation. The entries labelled Eigenvalues show the modulus of the largest three eigenvalues in the companion form of the VAR. Standard errors are calculated from Monte Carlo simulations based on the VAR estimates with 10000 replications.

Panel II reports variance decompositions at the one and three hour horizons.[13] Each row shows the percentage contribution by the source variable to the variance of the variable listed at the top of each column. The table also reports Monte Carlo standard errors based on 10,000 replications.

Several aspects of the results stand out. First, the variance of NT_t can be almost entirely accounted for by the innovations to trade and relative quote intensity at both the one and three hour

[13]These decompositions and the impulse response functions shown below are based on the variable ordering of $\{NT_t, NQ_t/NT_t, \sigma(T)_t\}$. The variance decompositions results are completely robust to this choice.

horizons. Heteroskedasticity in transaction prices plays no significant role. Second, trade intensity only makes a significant contribution to the variance of NQ_t/NT_t at the one hour horizon. Third, relative quote intensity contributes most to the variability in $\sigma(T)_t$ at both horizons, although the estimates are somewhat imprecise.

The impulse response patterns shown in Figure 7 provide complementary evidence on the dynamic relationship between these variables. [The dashed lines in the graphs show the 95% confidence band calculated by Monte Carlo simulation.] The left hand column shows the effects of an innovation in trade intensity. Here we see that a one standard deviation shock has a significant positive effect on trade intensity for approximately two hours. Unexpected changes in trading intensity appear to be fairly long lived. Trade innovations also have a persistent effect on the relative quote intensity. The ratio of quotes to trades is pushed below its normal level for more than one hour. By contrast, trade innovations have almost no significant effect on the volatility of transactions prices.

The middle column shows how unexpected changes in quote behavior impact upon transactions. An increase in the number of quotes [relative to the number of transactions] does lead to rise in transactions intensity that lasts for approximately one-and-a half hours, although the size of the rise is small. Quote shocks do not have any significant effect on the volatility of transactions prices. Thus, there is no support in these results for the idea that unexpected changes in trade or quote intensity significantly affect the subsequent volatility of transactions prices. Volatility shocks, by contrast, do affect subsequent trade and relative quote intensity. As the right hand column shows, a positive volatility shock significantly increases the subsequent number of transactions, and reduces the subsequent ratio of quotes to transactions. However, both effects are quite small. Finally, it is worth noting that volatility shocks are not very persistent. All the significant effects of a volatility shock are dissipated in 40 minutes.[14]

These findings have several noteworthy implications. First, variations in quote intensity do not appear to be a very accurate measure of trading intensity at high frequencies. In fact, the results suggest that quote intensity will understate trade intensity during periods of unusual market activity. Second, there is only week evidence that changes in quote behavior affect transactions. These findings would be hard to explain if quotes were firm rather than indicative. As it is, they once again served to underline that quote-making decisions are distinct form trading decisions in the FX market. The other noteworthy implication concerns the volatility of transactions prices. The volatility measure used here approximates the cross-sectional dispersion of transactions prices across the market. Since this measure is not seen directly by any trader, it is not surprising that shocks to the dispersion of transactions prices do not have a persistent effects. At the same time,

[14] These findings are robust to the ordering of variables. That is to say, the general impulse response patterns are not affected beyond that the first lag, and the standard error bounds indicate the same degree of statistical significance.

a greater dispersion of ¯rm quotes will generally raise the probability of a transaction taking place provided that the reservation price of the caller remains constant. This is likely because the increase dispersion of quotes across the market is unobserved. Thus, a higher dispersion of quotes will raise both the number and volatility of transactions prices consistent with results above.

5.2 Trades, Quote Volatility and Quote Spreads

The second VAR examines the relationship between trade intensity, quote volatility and the quote spread. Quote volatility is measured by the standard deviation in ask quotes during each ¯ve minute period, and the spread by the ¯ve minute average of the bid-ask spread. As above, the VAR was estimated with 6 lags and dummy variables to account for non-continuous observations.

Table 4: VAR results for Trade, Quote Volatiltiy and Spreads

I Granger Tests							
VAR Equation		NT_t		¾(Q)$_t$		SPR_t	
	NT_t	<0 001		0 001		0 011	
	¾(Q)$_t$	<0 001		<0 001		0 648	
	SPR_t	<0 124		0 027		<0 001	
II Variance Decompostions							
horizon	source	%	(Std. Err.)	%	(Std. Err.)	%	(Std. Err.)
1 hour	NT_t	99 957	(1 948)	77149	(14 507)	17411	(9170)
	¾(Q)$_t$	0 043	(1 467)	22 673	(14 476)	4 130	(6 334)
	SPR_t	<0 001	(1 279)	0 178	(2 629)	78 459	(10 530)
3 hours	NT_t	98 .704	(9 888)	99 .798	(10 216)	82 543	(18 001)
	¾(Q)$_t$	1 152	(5 078)	0 201	(4 178)	3 177	(7069)
	SPR_t	0 145	(8 510)	0 001	(8 983)	14 279	(15 .730)
III: Diagnostics							
Q_3		0 989		0 983		0 994	
Q_6		0 .783		0 .461		0 866	
3 Largest Eigenvalues		0 854		0 835		0 .790	

Notes: NT_t is the number of trades, ¾(T)$_t$ is the standard deviation of ask quote prices, and SPR_t is the mean bid ask quote spread, each calculated over a ¯ve minute period. The variables are ordered in the VAR as fNT_t;¾(T)$_t$;SPR_tg: For other estimation details, see Table 3.

Table 4 presents summary estimation results. Panel I shows that there is bi-directional Granger Causation between transaction intensity and quote volatility. This contrasts with the ¯ndings in Table 3 where transactions volatility did not appear to Granger Cause transactions. The quote spread appears exogenous with respect to quote volatility but is Granger Caused by transactions.

The variance decompositions reported in Panel II help to further clarify the relationship between the variables. It is immediately clear from these statistics the innovations in transactions play a far larger role in the behavior of quotes than was suggested by the causality tests. Innovations in transaction intensity are the major source of the variance in all three variables at the three hour horizon. Moreover, shocks to quote volatility explain little of the variance in either transactions or the spread. As above, these results are robust to the ordering of the variables.

Figure 8 plots the impulse response functions. The left hand column shows that a positive transaction shock leads to a persistence increase in quote volatility and a lowering of the quote spread. Interestingly, the latter effect occurs with a delay of approximately 30 minutes. The center column shows that while shocks to quote volatility temporarily raise the number of transactions, they do not significantly affect the spread. This difference in the behavior of the spread is robust to the ordering of the variables in the VAR. The right hand column shows that shocks to the quote spread have no significant affect on the other variables even though they are quite persistent.

These results point to an important distinction. Throughout I have stressed the quotes examined here are indicative rather than firm. However, this does not mean that quotes are unresponsive to developments in the market. Although the posting of quotes appears to be a distinct activity, the results above indicate that quote-making decisions are not made independently of trading activity. Rather, the response of quote volatility and the spread to transaction shocks appears consistent a competitive model for quotes in which banks compete to attract customers. Such competition insures that quotes respond to market developments even though they are only indicative. The results in Figures 7 and 8 provide the basis for examining this aspect of the FX market more fully.

5.3 Excess Demand, Transactions and Quote Prices

The final VAR examines the relation between excess demand, transactions prices and quotes. In the analysis above we saw that hourly changes in transactions prices were correlated with a cumulative measure of the excess demand for dollars. Here I examine the empirical foundation of this finding more closely. The VAR contains ED_t; the difference between the number of dollar purchases and sales during each five minute interval, ΔP^A_t; the change in the log purchase (ask) price over the interval, and ∇P_t; the difference between the logs of the last ask quote and the last purchase price during the interval. As above, the VAR includes 6 lags and dummy variables.

Panel I of Table 5 reports the marginal significance levels from the Granger Causality tests. We can reject the null of no Granger Causality from any of the variables in the case of both the excess demand and price change equations at the one percent level. By contrast, there is little evidence that any variable predicts the difference between quote and transactions prices, ∇P_t: In Panel II, the statistics indicate that the primary source of variation in all three variables are transaction

Table 5 : VAR results for Excess Demand, Prices and Quotes

I Granger Causality Tests			
VAR Equation	ED_t	¢ P_t^A	rP_t
ED_t	< 0 001	< 0 001	0 210
¢ P_t^A	< 0 001	< 0 001	0 474
rP_t	< 0 001	< 0 001	0 242

II Variance Decompostions							
horizon	source	%	(Std. Err.)	%	(Std. Err.)	%	(Std. Err.)
1 hour	ED_t	31 112	(10 670)	13 615	(27 599)	55 372	(26 210)
	¢ P_t^A	58 154	(13 508)	75 221	(30 341)	29 047	(27 985)
	rP_t	10 .734	(9 635)	11 164	(25 114)	15 581	(24 055)
3 hours	ED_t	24 259	(9 249)	24 324	(11 179)	24 275	(9 973)
	¢ P_t^A	64 163	(13 653)	64 104	(15 233)	64 127	(14 324)
	rP_t	11 082	(11 082)	11 572	(11 934)	11 597	(11 557)

III: Diagnostics			
Q_3	0 972	0 999	0 999
Q_6	0 418	0 999	0 999
3 Largest Eigenvalues	0 .713	0 579	0 579

Notes: ED_t is the di®erence between the number of dollar purchases and sales, ¢ P_t^A is the change in the log purchase (ask) price, and rP_t is the di®erence between the logs of the last ask quote and the last purchase price, each calculated over a ¯ve minute interval. The variables are ordered in the VAR as fED_t; ¢ P_t^A; rP_tg: For other estimation details, see Table 3 .

price shocks at the three hour horizon. Over the shorter horizon, excess demand shocks contribute more to the variance of ED_t and rP_t: As above, these results are robust to the variable ordering.

The impulse response patterns and associated standard errors derived from the VAR estimates are shown in Figure 9 . The impact of excess demand shocks are plotted in the left hand column. Here we see that while the major impact of these shocks on excess demand occurs within the period [i.e., ¯ve minutes], small but signi¯cant e®ects persist for approximately 40 minutes. Positive excess demand shocks also signi¯cantly reduce ¢ P_t^A and raise rP_t for two periods. Notice that the negative correlation between excess demand and transactions prices implied by this response pattern is the opposite of the positive correlation in hourly data documented in Table 2 .

Further insight into the origins of the hourly correlation comes from the center column of Figure 9 where the impact of shocks to transactions prices are shown. Here we see that these shocks have a positive and then negative impact on ¢ P_t^A consistent with a one period a®ect on the level of prices, P_t^A. Once again the response of quotes appear to lagged behind prices so that ∇P_t falls temporarily. The impact of price shocks on excess demand is much more persistent. As the upper plot shows, positive price shocks significantly raise excess demand for up to one hour. Thus, price shocks have a more persistent effect on excess demand than demand shocks, a fact that is reflected in the variance decompositions above. There we saw that price shocks are the major contributor to the variance of excess demand at both the one and three hour horizons. This response pattern also appears more consistent with the positive correlation between prices and excess demand in the hourly data.

The right hand column of the figure plots the impact of innovations to ∇P_t. Since these innovations are uncorrelated with price shocks [by construction], shocks to ∇P_t provide a measure of the unexpected changes in quotes. As the figure shows, these quote shocks are not very persistent - there is no significant effect on ∇P_t beyond one period. The impact on transaction prices and excess demand is more interesting. The response of ΔP_t^A indicates that transaction prices rise permanently in the period following a quote shock. Positive quote shocks also lead to a significant rise in excess demand that last for approximately 30 minutes. Neither of these findings squares with the view that quotes are simply noisy measures of transactions prices. If that were the case, why would the "noise" measured by ∇P_t affect trading activity in the way indicated by the impulse response patterns? A more plausible explanation appears to be that quotes respond to market-relevant news, such as customer orders.

These findings merit careful analysis. They suggest a complex dynamic interaction between transactions and prices that needs to be "explained" by theoretical models of the trading process. In particular, such a model will need to explore why the response of excess demand and prices can differ so much according to the origin of the shock. There also appears to be a complex relationship between transaction quantities, prices, and indicative quotes. The results above in Figure 8 indicate that quote-making decisions are not made independently of trading activity. Here we see evidence of trading decisions being affected by quotes. As yet, we have no theoretical understanding of these relationships.

6 Conclusion

By its very nature, a paper such as this raises many more questions than it answers. Since my aim here was to provide a summary analysis of a new data set, I have not followed many of the research avenues opened up by the results. Consequently, a lot more work needs to be done before we will have learned all we can about the microstructure of the FX market from these data. Nevertheless, I think is safe to draw the following points from the analysis above.

First, trade activity within the interbank FX market is distinct from the posting of indicative quotes. Trading and quote-making decisions are linked, but the links are complicated and poorly understood. Second, microstructure e®ects appear important in the determination of exchange rates at high frequencies. There is also some evidence that these e®ects are important over the longer time spans relevant in international macroeconomics. Could this mean that a better understanding of the microstructure of the FX market would aid in resolving the big puzzles in international macroeconomics? While it is obviously much too soon to say, it is an intriguing possibility.

References

Andersen, T. and T. Bollerslev, 1994, "Intraday seasonality and volatility persistence in foreign exchange and equity markets", Kellogg Graduate School of Management, Northwestern University, Working Paper # 193.

Baillie, R. and T. Bollerslev, 1991, "Intra day and inter market volatility in foreign exchange rates", Review of Economic Studies, 58, 565-585.

Bessembinder, J., 1994, "Bid-ask spreads in the interbank foreign exchange market", Journal of Financial Economics, 35, 317-348.

Bollerslev, T. and I. Domowitz, 1993, "Trading patterns and prices in the interbank foreign exchange market", Journal of Finance, 48, 1421-1443.

Bollerslev, T. and M. Melvin, 1994, "Bid-ask spreads and volatility in the foreign exchange market", Journal of International Economics, 36, 355-372.

Dacorogna, M., Muller, U., Nagler, R., Olsen, R. and O. Pictet, 1993, "A geographical model for the daily and weekly seasonal volatility in the FX market", Journal of International Money and Finance, 12, 413-438.

Demos, A. and C. Goodhart, 1992, "The interaction between the frequency of market quotations, spread, and volatility in the foreign exchange market", LSE Financial Markets Group Discussion Paper, No. 152, October.

Engle, R., Ito, T. and W-L. Lin, 1990, "Meteor showers or heat waves? Heteroskedastic intradaily volatility in the foreign exchange market", Econometrica, 58, 525-542.

Glassman (1987), "Exchange rate risk and transactions costs: Evidence from bid-ask spreads," Journal of International Money and Finance

Goodhart, C. and M. Giugale, 1993, "From hour to hour in the foreign exchange market", The Manchester School, 61, 1-34.

Goodhart, C., Ito, T. and R. Payne, 1996, "One Day in June 1993: A study of the working of Reuters 2000-2 electronic foreign exchange trading system," in Frankel, Galie and Giovannini (eds.) The Microstructure of Foreign Exchange Markets, Chicago

Guillaume, D., Dacorogna, M., Dave, R., Muller, U., Olsen, R. and O. Pictet, 1994, "From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets", A discussion paper by the O&A research group, Zurich, October.

Hartmann, P., 1997, "Do Reuters spreads reflect currencies' differences in global trading activity?", FMG Disscusion Paper No. 265, LSE.

Ito, T., R. Lyons, and M. Melvin, 1997, "Is there private information in the FX market? The Tokyo experiment", Journal of Finance, forthcoming.

Ito, T., Engle, R. and W-L. Lin, 1992, "Where does the meteor shower come from?", Journal of International Economics, 32, 221-240.

Lyons, R., 1996 "Foreign exchange volume: sound and fury signifying nothing?" in Frankel, Galie and Giovannini (eds.) The Microstructure of Foreign Exchange Markets, Chicago

Lyons, R., 1995 "Tests of microstructure hypotheses in the foreign exchange market," Journal of Financial Economics

Lyons, R., 1997 "A simultaneous trade model of the foreign exchange hot potato", Journal of International Economics, May, 275-298.

Muller, U., Dacorogna, M., Olsen, R., Pictet, O., Schwartz, M., and C. Morgenegg, 1990, "A Statistical study of foreign exchange rates, empirical evidence of a price scaling law, and intraday analysis", Journal of Banking and Finance, 14, 1189-1208

Muller, U., and Sgier, R., 1992, "Statistical analysis of intra-day bid-ask spreads in the foreign exchange market", Internal document UAM 04-10, Olsen & Associates, Zurich.

O Hara, M. 1994, Market Microstructrue Theory, Blackwell, Oxford.

Payne, R., 1996, "Announcement effects and seasonality in the intra-day foreign exchange market", LSE Financial Markets Group Discussion Paper, No. 238, March

Wei, S-J., 1994, "Anticipations of foreign exchange volatility and bid-ask spreads", NBER Working paper No. 4737, May.

Yao (1996) "Market making the interbank foreign exchange market," mimeo, Stern School of Business, NYU.

Figure 1

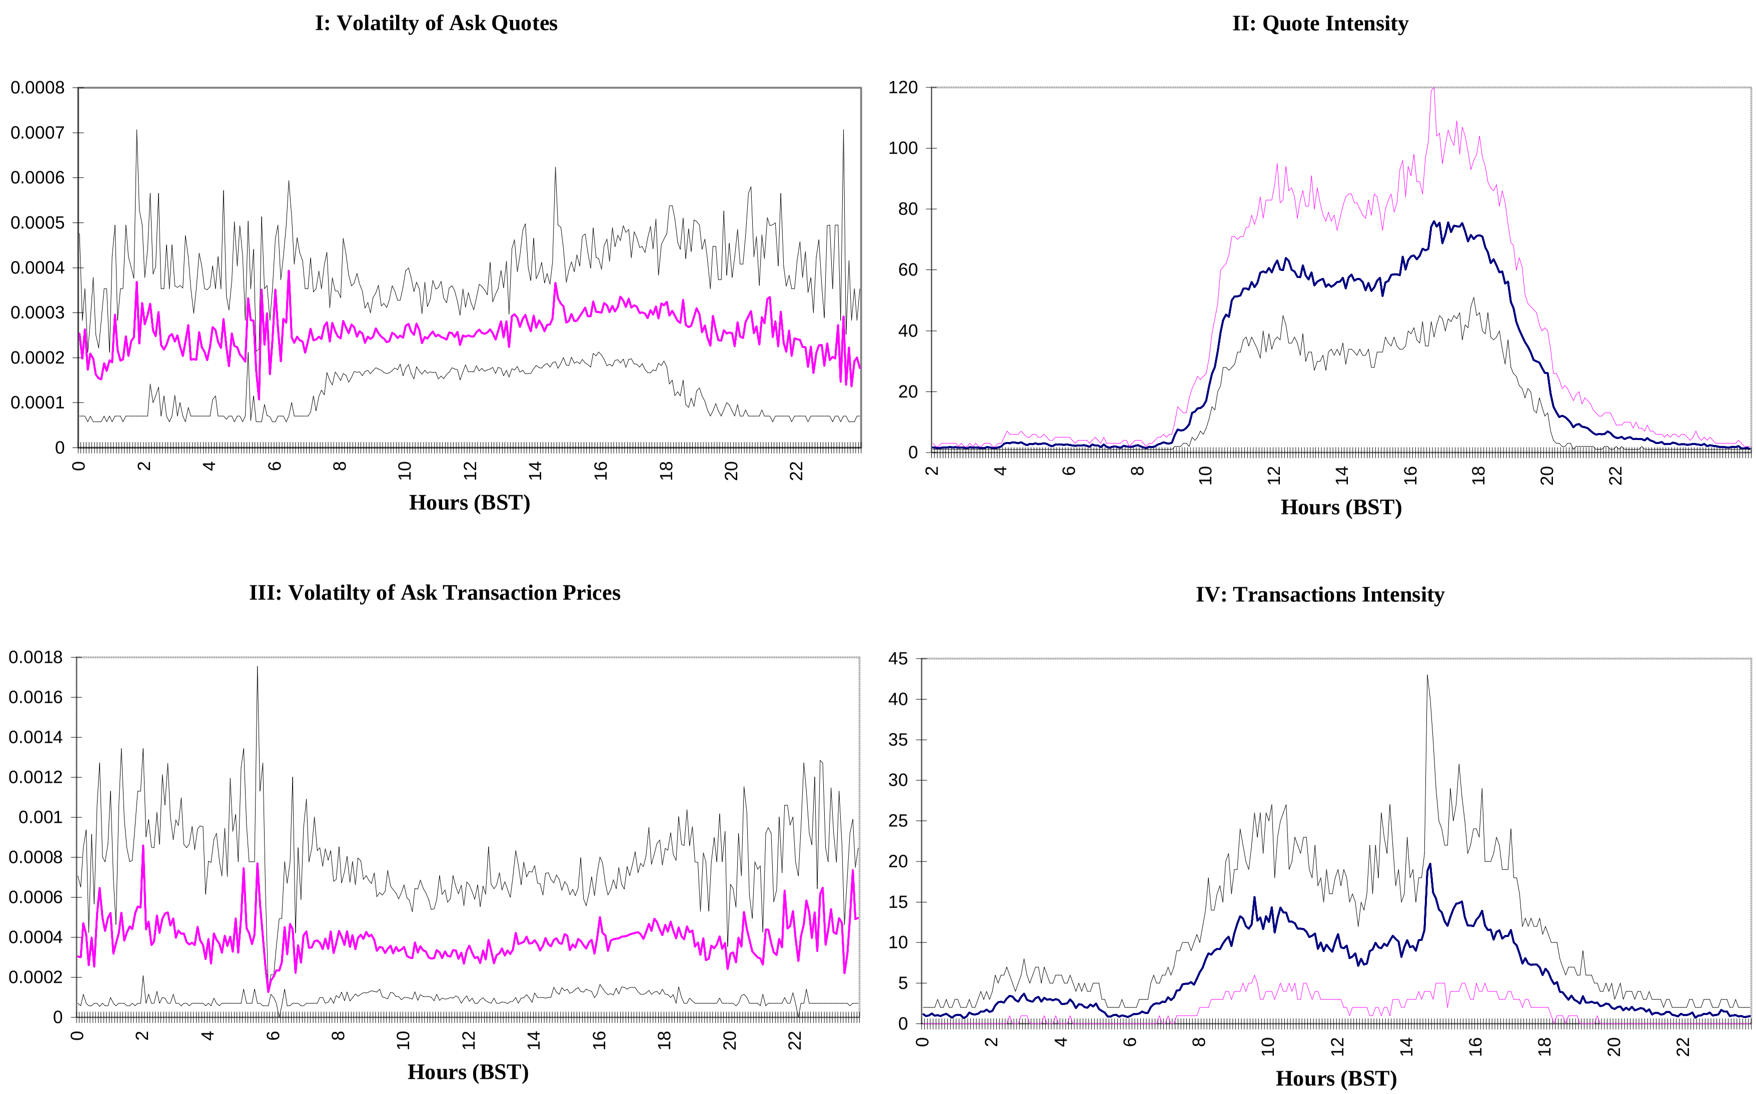
I: Volatilty of Ask Quotes
0.0008
0.0007
0.0006
0.0005
0.0004
0.0003
0.0002
0.0001
0
0 2 4 6 8 10 12 14 16 18 20 22
Hours (BST)
II: Quote Intensity
120
100
80
60
40
20
0
2 4 6 8 10 12 14 16 18 20 22
Hours (BST)
III: Volatilty of Ask Transaction Prices
0.0018
0.0016
0.0014
0.0012
0.001
0.0008
0.0006
0.0004
0.0002
0
0 2 4 6 8 10 12 14 16 18 20 22
Hours (BST)
IV: Transactions Intensity
45
40
35
30
25
20
15
10
5
0
0 2 4 6 8 10 12 14 16 18 20 22
Hours (BST)

Figure 2

I: Mean Quote Spread

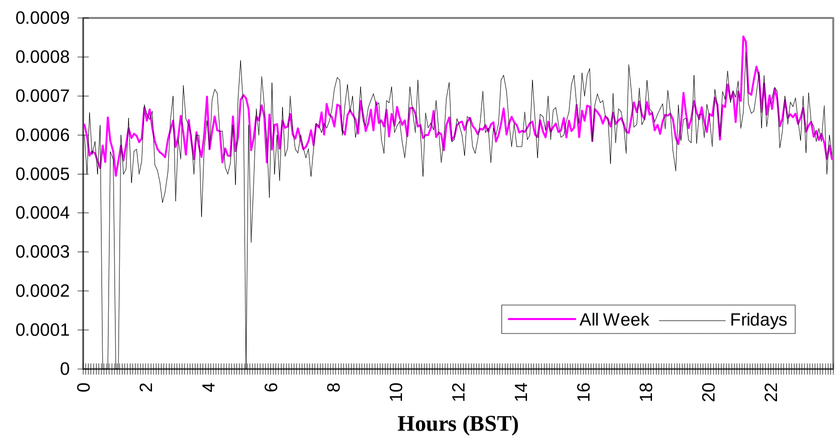

II: Volatility of Quote Spread

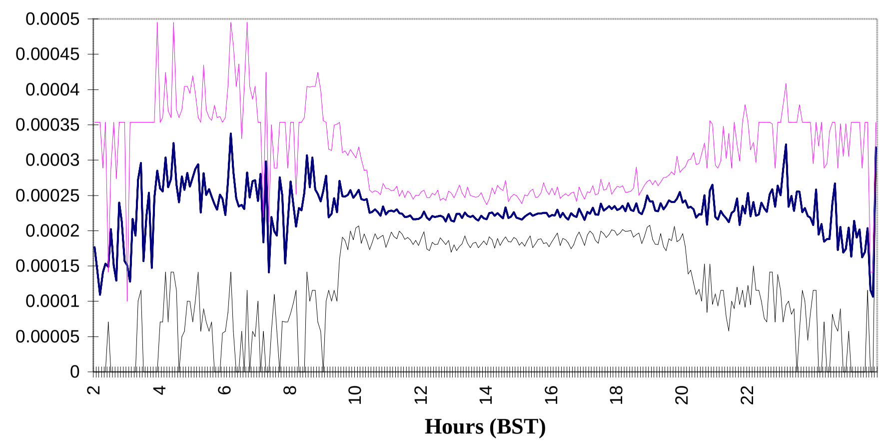

III: Ask Transaction Price Range

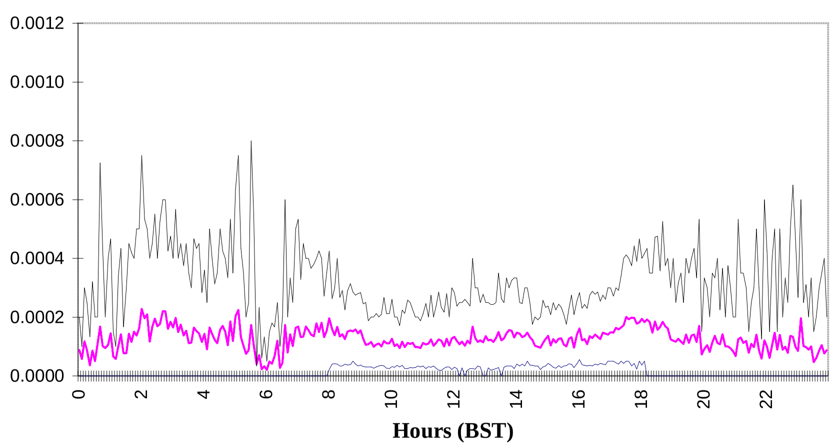

IV: Bid Transaction Price Range

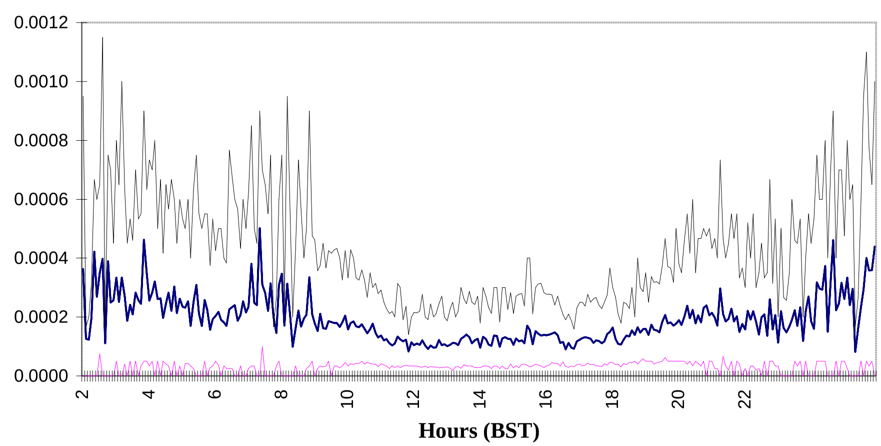

Figure 3

I:Volatility of Ask Quotes

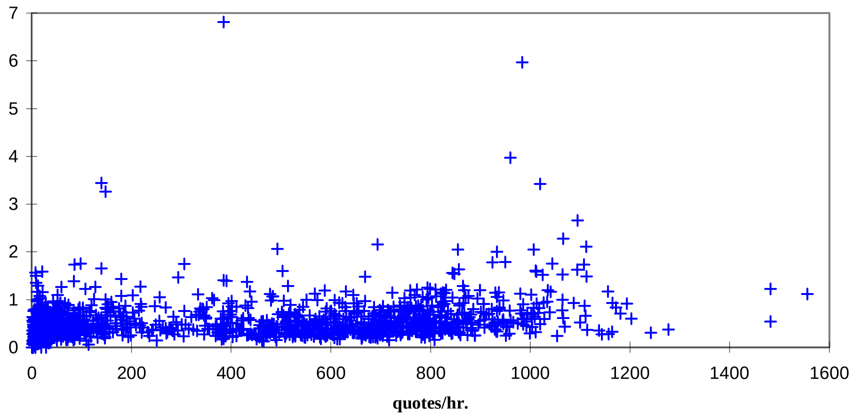

II:Volatility of Ask Transaction Prices

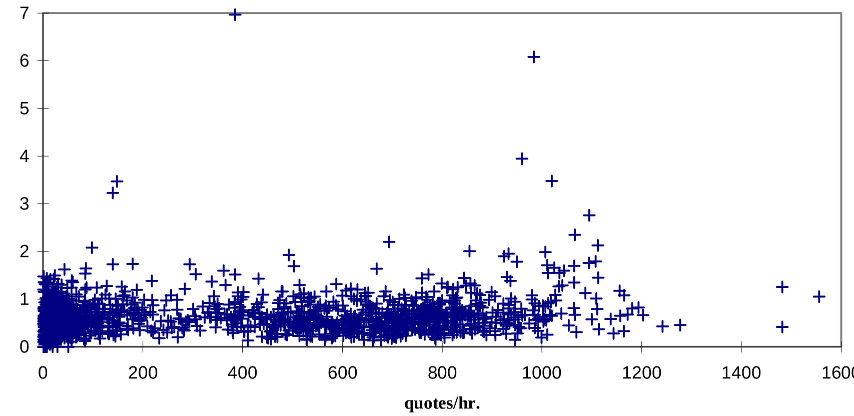

III:Trades per hour

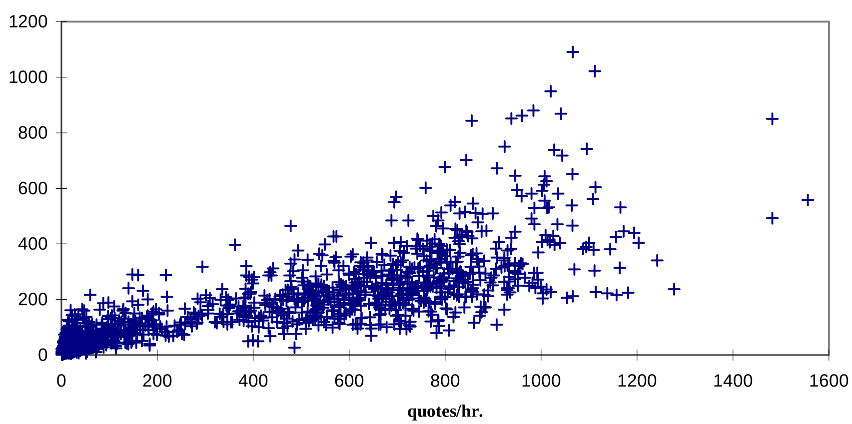

IV:Trade Volume ($Bn./hour)

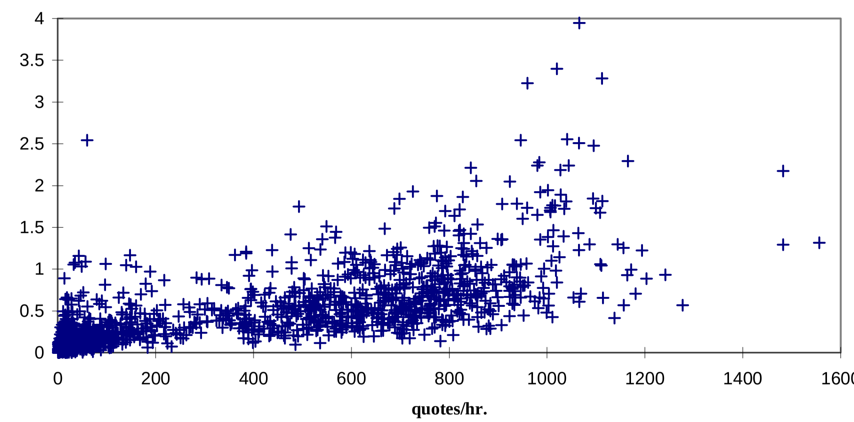

Figure 4

I:Volatility of Ask Quotes

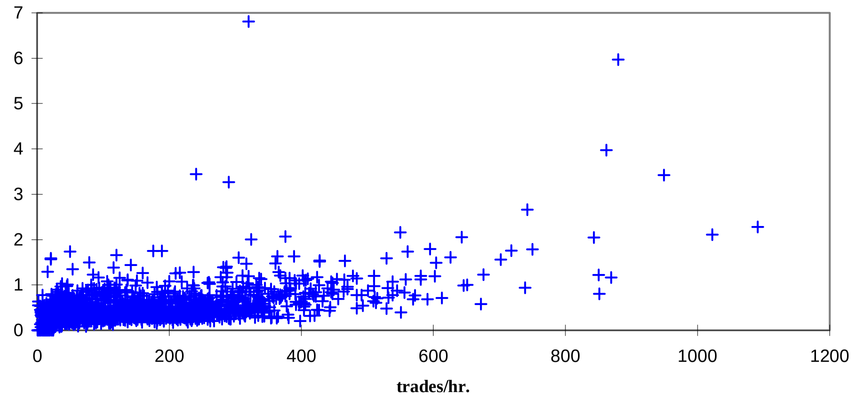

II:Volatility of Ask Transaction Prices

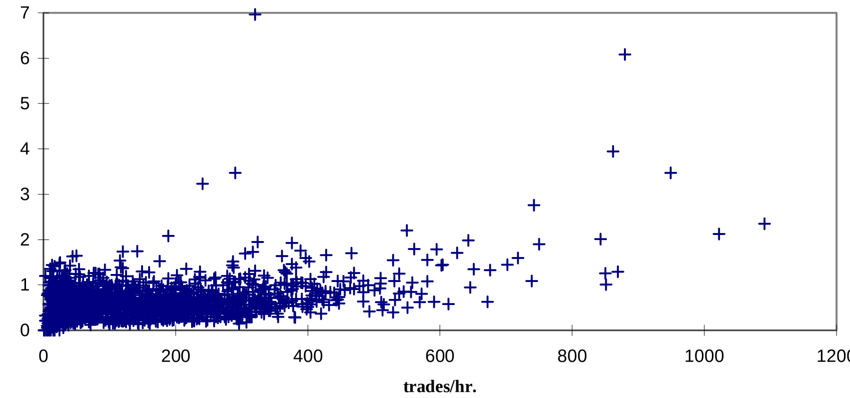

III:Mean Quote Spread

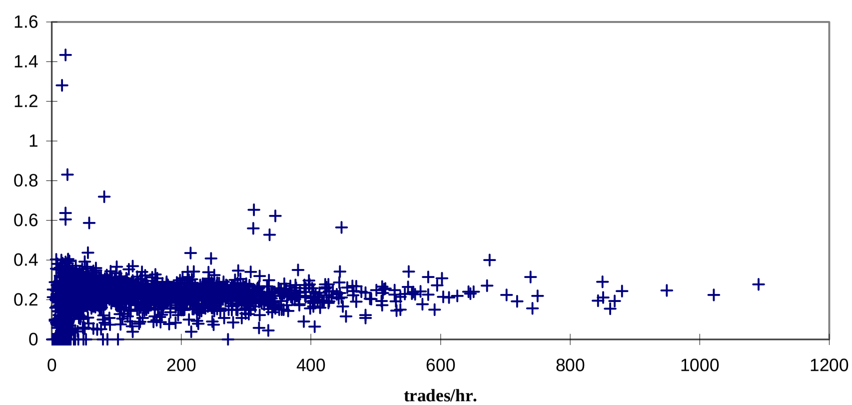

IV:Trade Volume ($Bn./hour)

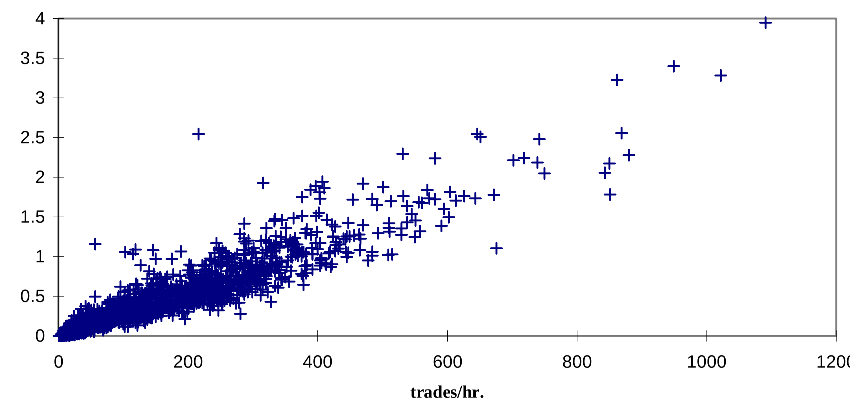

Figure 5

I:Volatility of Ask Quotes

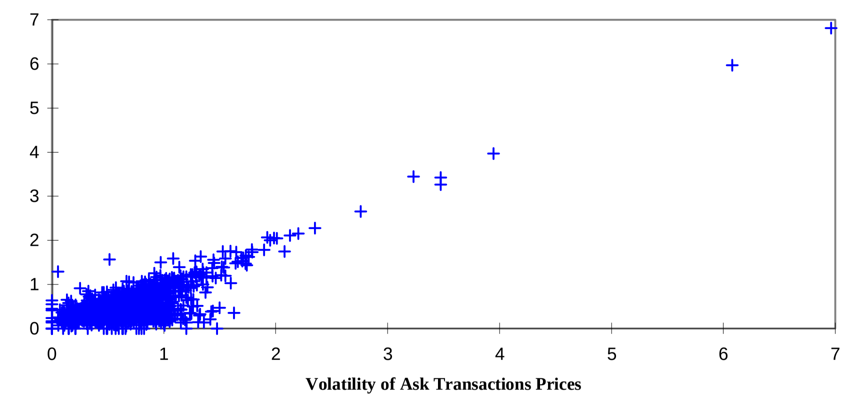

II:Range of Ask Quotes

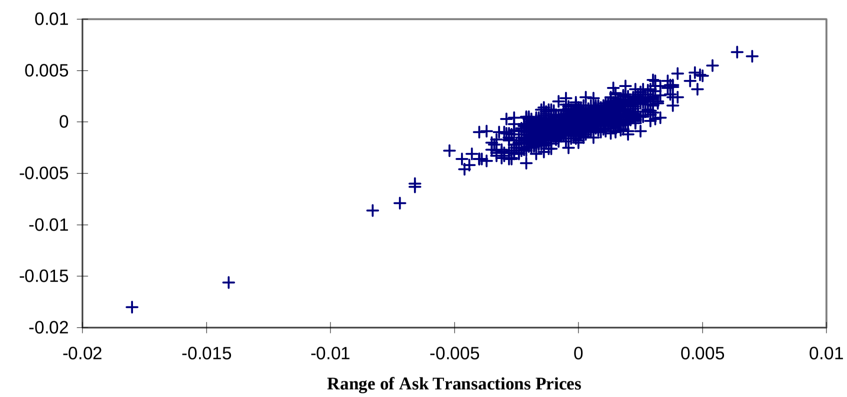

III:Mean Quote Spread

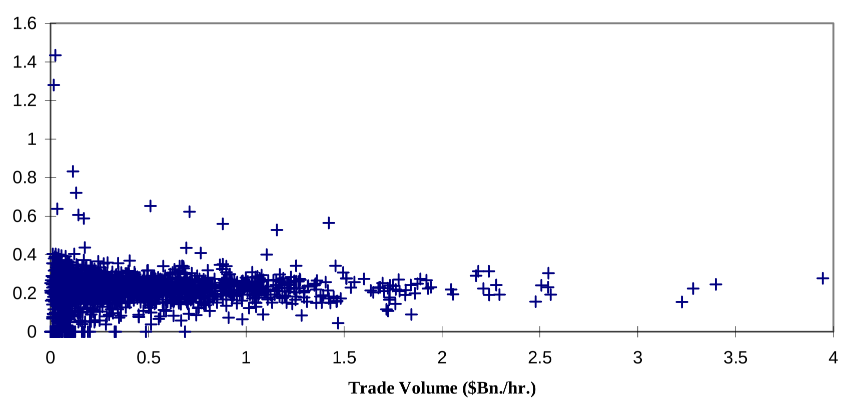

IV:Volatility of Ask Quotes

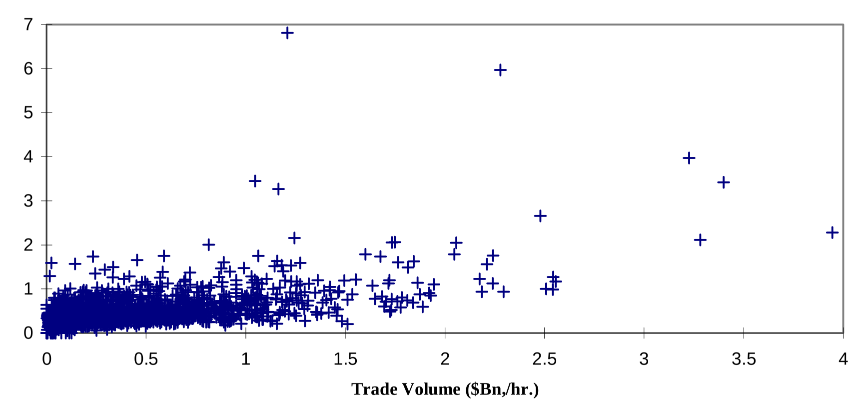

Figure 6

I:Daily Change in Ask Transactions Price

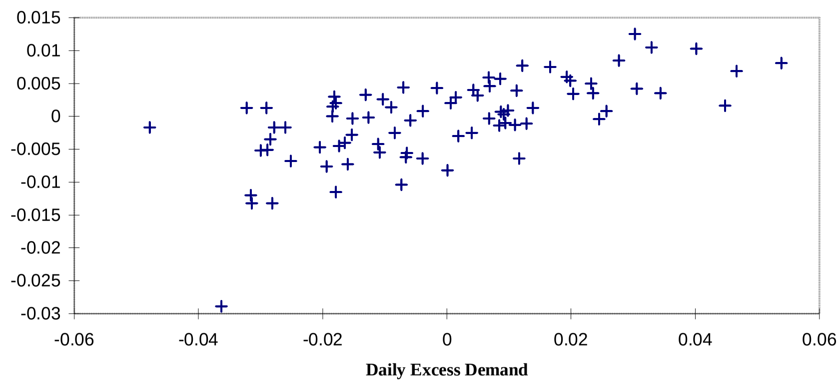

II:Cumulative Daily Trade Imbalance

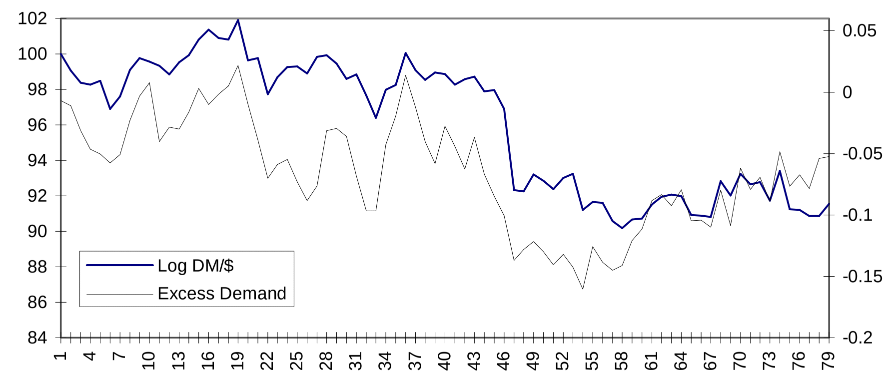

III:Cumulative Change in Hourly Ask Price

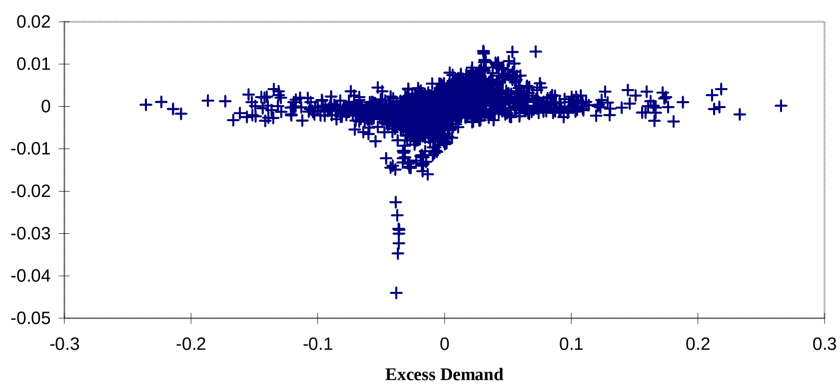

IV:Cumulative Change in Hourly Ask Price (PM)

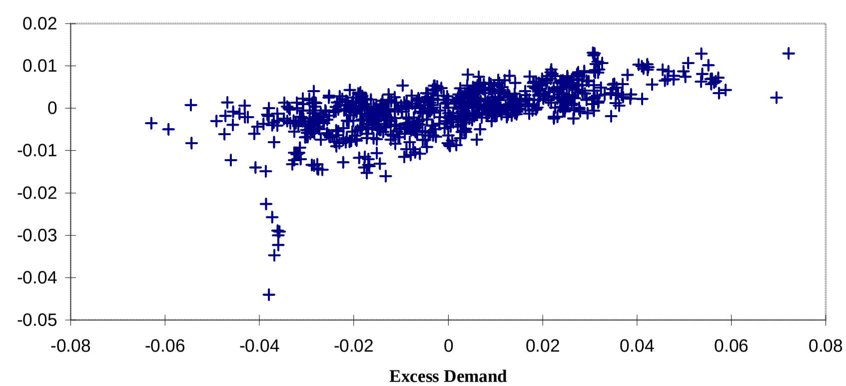

Figure 7

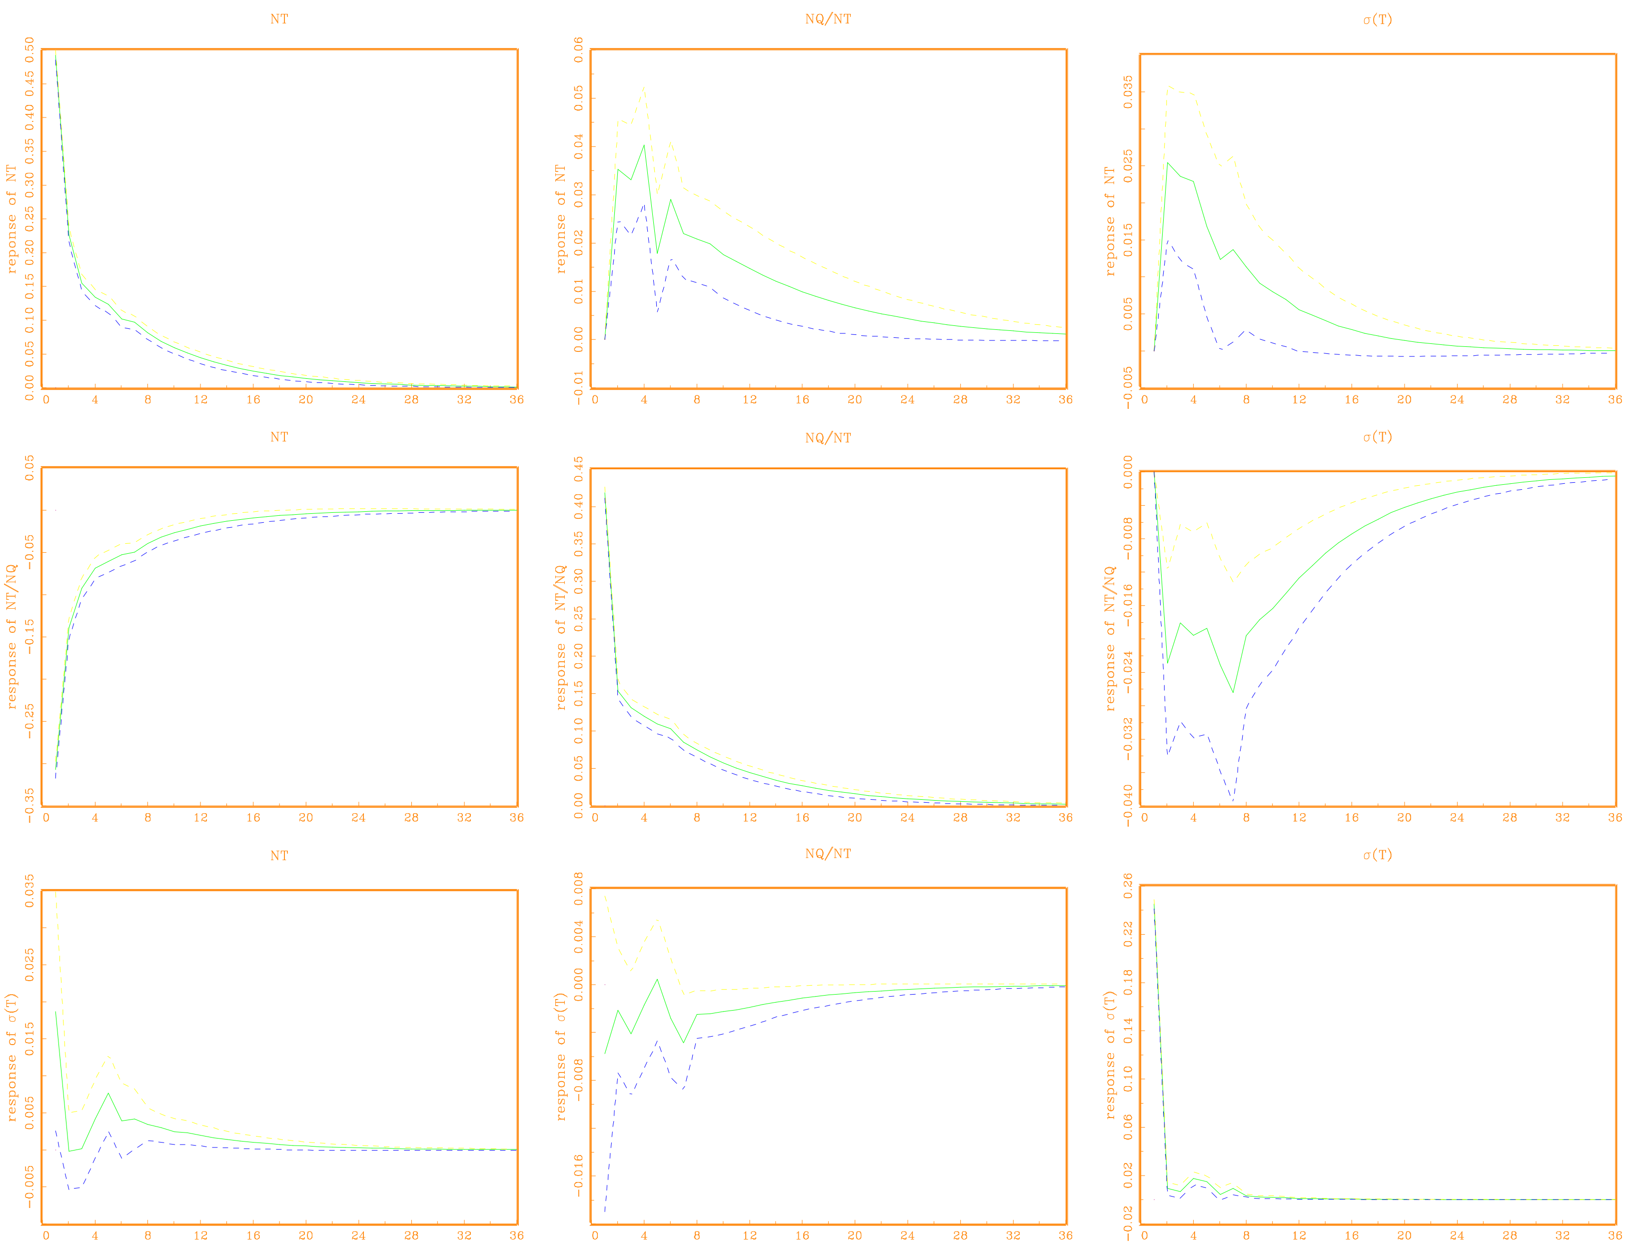
NT
NQ/NT
σ(T)
reponse of NT
response of NT/NQ
response of σ(T)
0 4 8 12 16 20 24 28 32 36

Figure 8

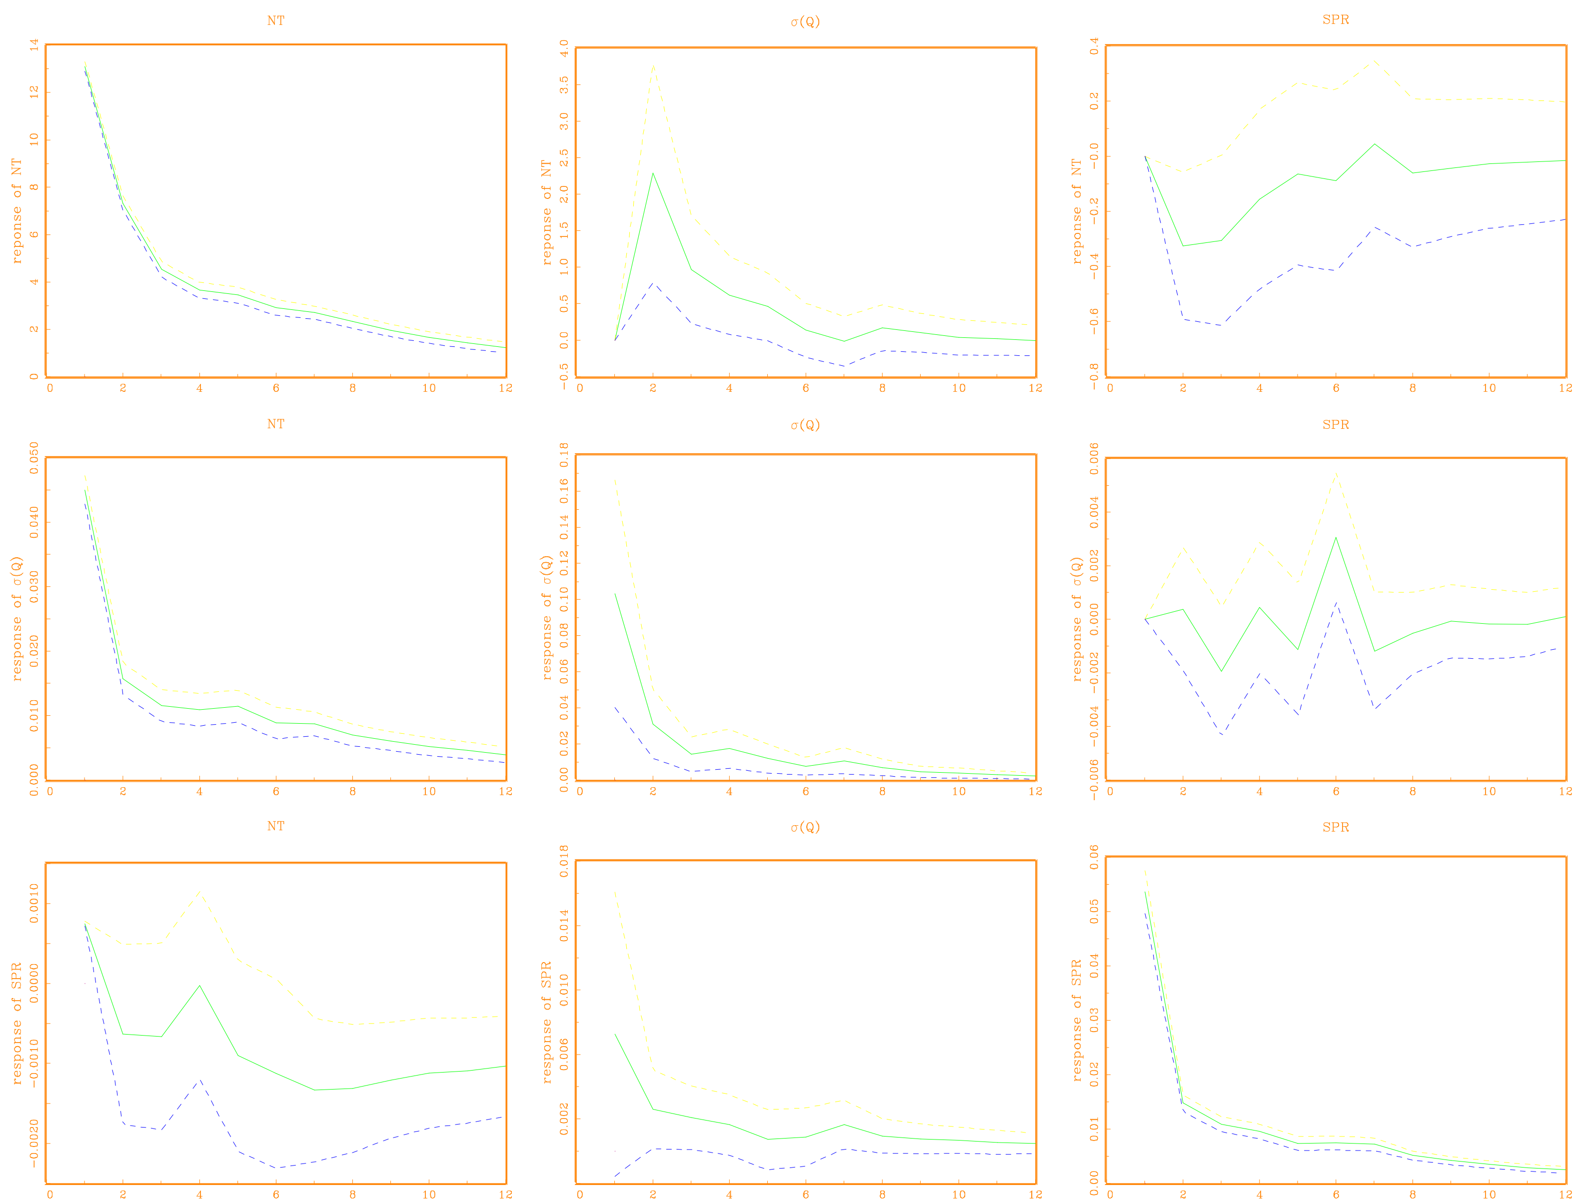

Figure 9

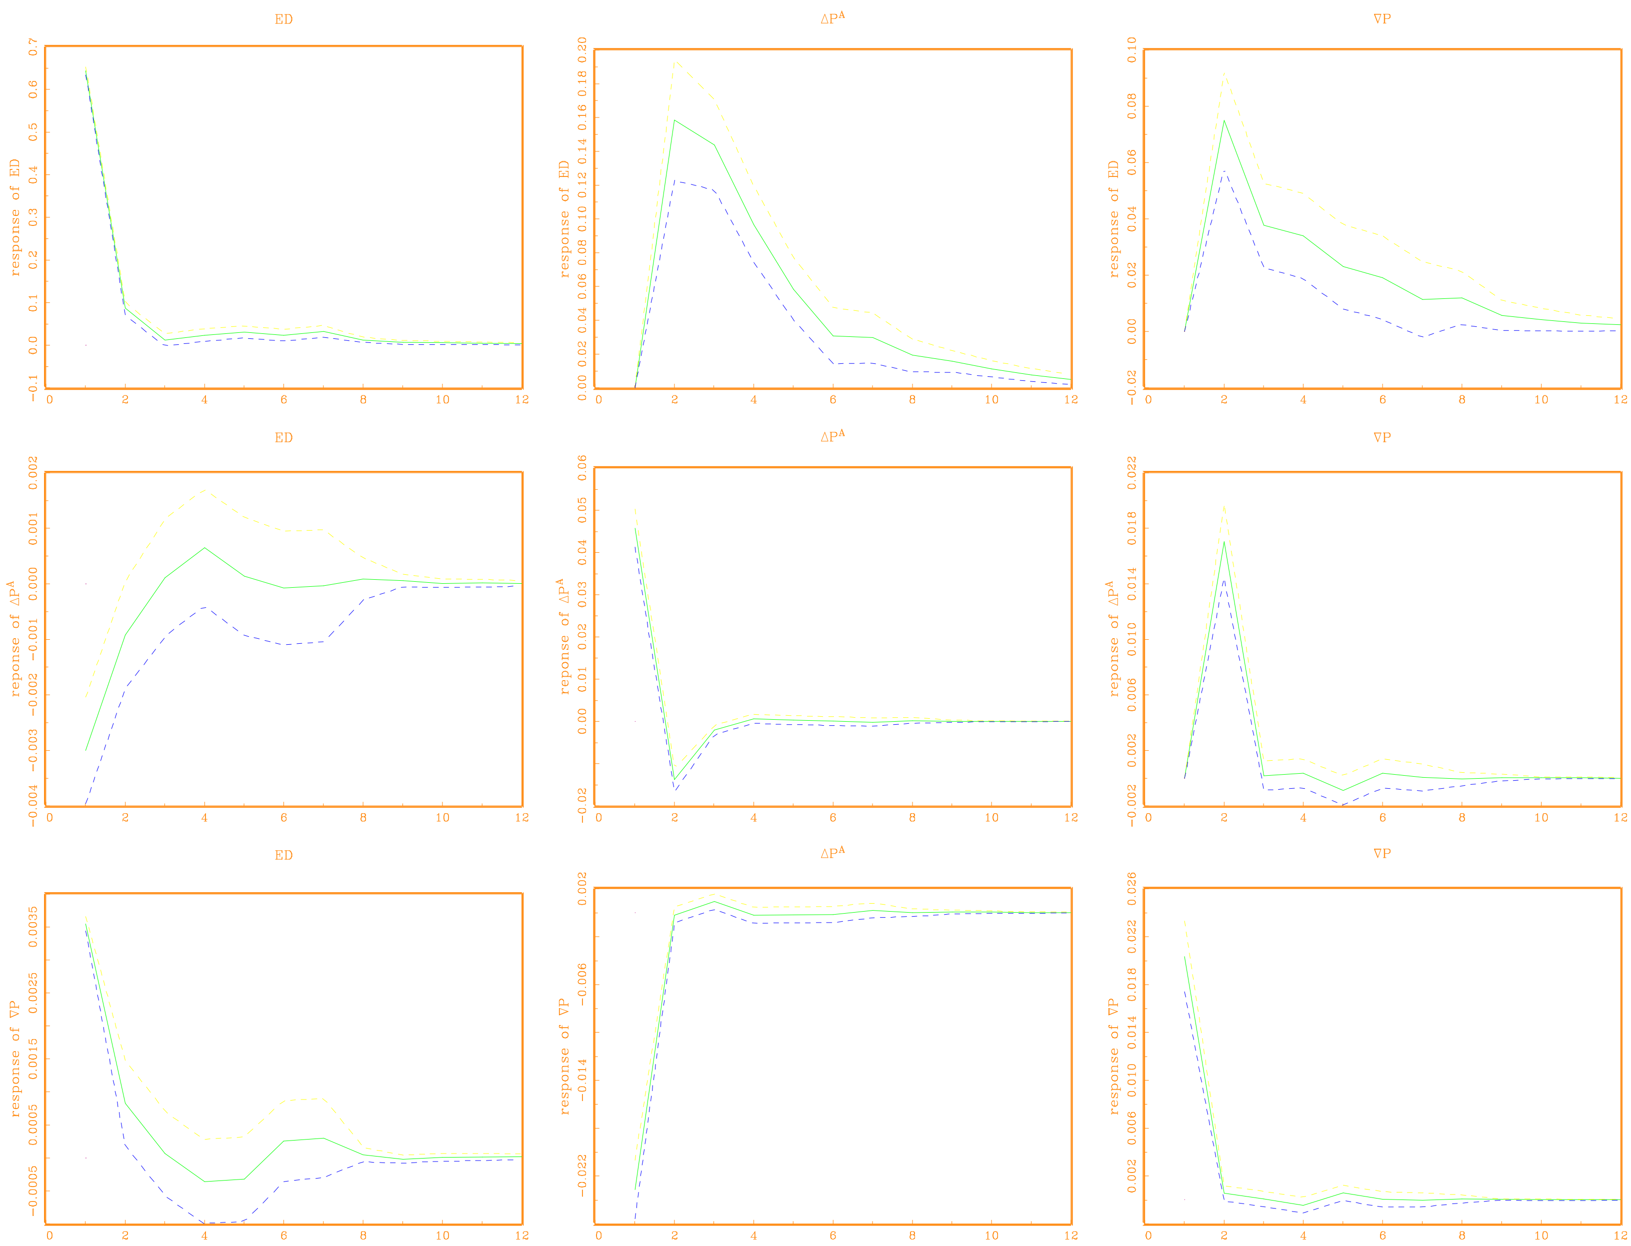

ED
ΔP^A
∇P
response of ED
response of ΔP^A
reponse of ΔP^A
response of ∇P