Informed Traders and Dealers in the FX Forward Market

Pierre Collin-Dufresne, Peter Hoffmann and Sebastian Vogel*

December 03, 2019

Abstract

There is strong heterogeneity in the permanent price impact of traders. Moreover, a trader's permanent price impact is persistent. A trade's ex-post permanent price impact is partially priced in dealers' markups, even when controlling for dealer-client fixed effects. This suggests that dealers are informed about the permanent price impact of their clients' trades. We present further evidence suggesting that dealers learn from their clients' order flow and use this knowledge when providing quotes. More informed customers are more likely to trade with informed dealers. We present a model explaining our empirical findings.

1 Introduction

The two-tiered market structure, where clients' trades are intermediated by dealers who can trade amongst each other in an interdealer market, remains prevalent in many OTC markets including fixed income, credit, and foreign exchange. In theory, such a structure may naturally arise if clients are differentially informed, and can signal their type to dealers who can price-discriminate (Seppi (1990), Lee and Wang (2019)). More recently, Glode and Opp (2016, 2019) argue that, in the presence of asymmetric information, intermediation chains may be needed to generate efficient trading behavior. A property of these intermediation chains is that trades are expected to occur between counterparties that are similarly informed. In practice however, it is less clear how these efficient intermediation chains can arise[1] and whether we can empirically observe that trading relationships are, at least to some extent, determined by informedness of the different traders.

*Pierre Collin-Dufresne is at EPFL and the Swiss Finance Institute (email: pierre.collin-dufresne@epfl.ch). Peter Hoffmann is at the European Central Bank (email: peter.hoffmann@ecb.int). Sebastian Vogel is at EPFL and the Swiss Finance Institute (email: sebastian.vogel@epfl.ch). We gratefully acknowledge the financial support of the European Central Bank and the European Systemic Risk Board through the EMIR Bridge Programme. The views expressed in this paper are those of the authors and do not necessarily represent the views of the institutions to which they are affiliated.

[1] Glode and Opp (2016, 2019) consider the case in which the trading network is held fixed. If this would not be the case, informed traders had incentives to choose less informed counterparties.

In this paper we use data on foreign exchange transactions made available through the European Market Infrastructure Regulation (EMIR) to shed new light on the functioning of one of the largest OTC markets: the euro-dollar forward exchange rate market. This market is largely two-tiered in that most client transactions occur with a limited set of dealers and there is a very active interdealer market. Since the data set contains all the individual transactions in the EU with information on trader identities, we can ask the following questions. Is there evidence that clients are differentially informed? Is there evidence that dealers are differentially informed? Are markups charged by dealers related to client and/or dealer informedness? Are the client-to-dealer (and dealer-to-dealer) trading networks affected by these differences in informedness?

To answer the first question, we measure price impact of individual clients' trades at different horizons (1-minute, 30-minutes, 1-day). A positive price impact implies that clients tend to buy (sell) from the dealer when benchmark rates increase (decrease) subsequent to their trade. Such situation should arise if clients are on average better informed than dealers about future exchange rate changes. This is the standard adverse selection mechanism presented in the traditional microstructure literature (e.g., Kyle (1985) and Glosten and Milgrom (1985)). Of course, it is perhaps less intuitive to think that clients have private information about future exchange rate fundamentals, which we typically think of as reflecting macro-economic risks. However, as we show in a simple theoretical model, clients information may pertain to their individual order flow, which may be correlated with total order flow, which in the short run may affect the change in exchange rates (see also Evans and Lyons (2002, 2005)). Alternatively, some investors may also be better at interpreting public news, say about macroeconomic fundamentals, and thus effectively also have private information about systematic sources of risk (e.g., Kim and Verrecchia (1991)). Empirically, we find that on average clients' price impact is highly statistically significant and positive at a 1-minute horizon. At a longer (1-day) horizon it remains highly statistically significant and positive for hedge funds on average. However, there is considerable cross-sectional and time-series variation across traders. When we look at the individual trader level, we find significant persistence in price impact. Breaking down the sample into subperiods, we find that traders that tend to have a higher price impact in the first subperiod tend to remain in the high-price impact group in subsequent periods. This suggests that some (groups of) traders are consistently better informed, in that their trades seem to, on average, correctly anticipate future exchange rate changes.

At some level, these findings are consistent with the original findings of Evans and Lyons (2002), who documented, using 4 months of data in 1996, that it was possible to predict future exchange rate changes based on aggregate interdealer order flow. Since, as we show in our theoretical model, one would expect the interdealer order flow to be driven by their clients' order flow, it seems natural to anticipate that there should be some information in at least some of the clients' trades. We confirm this intuition by extending the Evans and Lyons (2002) study to the individual dealer level. Specifically, we investigate whether the aggregated clients' order imbalance observed by each dealer allows them to predict future exchange rate movements. We find strong evidence of predictability at the dealer level. That is, individual dealers could earn significant Sharpe ratios from trading based on their clients aggregated order flow. However, there is also substantial cross-sectional variation across dealers. We label the dealers with the highest predictive client order flow, the 'informed dealers,' and we study the characteristics of these dealers and whether we see specific patterns in the client-dealer trading network. Interestingly, we find that dealer informedness is not isomorphic to the standard centrality measures such as connectedness. We find that more informed dealers indeed use their information when giving quotes to traders. Markups are generally higher, the higher the price impact of the trade. The more informed the dealer, the stronger is this effect. We also find that traders are more likely to trade with informed dealers if they are informed themselves. Relatedly, for all traders that are not HFT whose investment horizon is arguably very short, traders are more likely to trade with informed dealers if volatility, a proxy for adverse selection, is high.

To interpret our empirical findings we develop a simple model of a two-tiered OTC market in which dealers intermediate trades between their customers and subsequently hedge their inventory risk in an interdealer market. In this model, order flow is informative about future price changes. Moreover, some dealers may forecast future price changes better than other dealers. Those dealers not only incorporate some of their information in their markups, but are also more likely to further attract informed traders due to an adverse selection problem that the less informed dealers face: If an informed trader asks an uninformed dealer for a bid and an ask price, this dealer likely has difficulties to respond as the willingness of the informed trader to buy or sell at a given ask or bid price will mean bad news for the uninformed dealer. On the other hand, the informed dealer is able to correctly price an asset and is able to make a market for informed traders.

By focusing on the informedness of traders, we extend a growing literature on the network structure of OTC markets.[2] While Wang (2017) develops an inventory-based model of the OTC market structure and Sambalaibat (2018) develops a model in which dealers specialize on the trading frequency of their clients, our findings suggest that dealers also specialize based on the informedness of their customers. In the light of theoretical and empirical results of Babus and Kondor (2018) and Kondor and Pintér (2019) it is surprising that the more informed market participants have fewer counterparties than their less informed counterparts. However, this result is not unreasonable, since informed counterparties are especially vulnerable to infor-

[2]Characteristic for OTC markers is a core-periphery structure, see for instance Abad et al. (2016), Li and Schürhoff (2019) or Neklyudov et al. (2017).

mation leakage that arises from being in contact with many counterparties.[3] While previous research on a two-tiered OTC market pointed out the benefits of trading with connected dealers, who can provide more immediacy (Di Maggio et al. (2017) and Li and Schürhoff (2019)), we show in this paper that it can sometimes be beneficial to trade with a less connected dealer: In the presence of strong information asymmetries, only informed dealers may be willing to provide quotes to informed traders. However, these dealers generally have fewer counterparties.

Also Bjønnes et al. (2017) and Ranaldo and Somogyi (2018) study informed trading in the FX market. The advantage of the EMIR dataset compared to the datasets used in these studies lies in the availability of the traders' and the dealers' identities. Thus, we can not only document the presence of informed trading, but also study the persistence of informedness and characteristics as well as the trading behavior of informed traders with different dealers.

In an influential paper, Evans and Lyons (2002) show that aggregate orderflow predicts future price changes in the FX market. Menkhoff et al. (2017) show that the order flows from different subsets of traders have different forecasting abilities. We connect to this strand of literature by showing that the order flow of different dealers has different forecasting abilities. Moreover, the informedness measures we obtain for traders and dealers allow us to study the traders' dealer choice problem.

2 Data and Summary Statistics

We use three different databases: The EMIR database contains information on derivatives transactions in which at least one counterparty is located in the EU. We use the full database to which the ESRB and ESMA have unique access. We focus on the FX forward market for the following reasons. First, it is one of the largest derivatives markets. As shown in Abad et al. (2016), the FX forward market is the second largest derivatives market in the EU. While the interest rate swaps (IRS) market is still larger in terms of notional volume, Abad et al. (2016) show that approximately 85% of notional volume of IRSs is being traded among G16 dealers and other banks. Nonfinancial firms only generate less than 1% of the notional volume traded in the IRS market. Traders in the FX forward market are more diverse. Less than 70% of notional volume is generated among G16 dealers and other banks.

The other two databases we use are the ORBIS database, which contains information on the different

[3] Hendershott and Madhavan (2015) argue that information leakage is an important concern when evaluating whether to contact many dealers via an RFQ trading protocol or a single dealer in the voice market. Hagströmer and Menkveld (2019) measure information flow between dealers in the FX market and Liu et al. (2018) model the information leakage in a two-tiered OTC market, showing that informed traders may benefit from limiting the number of contacted dealers.

types of traders, and the Thomson Reuters Tick History (TRTH) from which we derive benchmark prices for the forward contracts. The following subsections describe the data in more detail.

2.1 EMIR and ORBIS data

We use the EMIR activity report and focusing on the last message submitted for each trade. Moreover, we look at the period from May 2018 to April 2019 and restrict our attention to transactions that happen between Monday and Friday as well as between 8am and 8pm UTC and exclude transactions with no reported price rate, or markups with an absolute value of more than 5 %. How we determine the markup for each trade is explained further below in this section. As in Abad et al. (2016) and Hau et al. (2019) we use the ORBIS database to assign a type to each trader. Possible types are FUND, BANK, G16, INSURANCE & PENSION, NON-FINANCIAL, CENTRAL BANK and EMPTY. Firms not covered by the ORBIS dataset are also classified as EMPTY. The EMIR database reports only the legal entities that were involved in a transaction. Many firms, especially the G16 dealers use many legal entities. The ORBIS dataset allows us to associate each legal entity with its parent company.

Table 1: Averages of firm characteristics for D2C market. This table shows the number of firms for each type in the sample, shows how large the notional volume in EUR per trade involving these firms is and looks at the trader characteristics for trades involving these firms. The trader characteristics include the number of monthly counterparties and average monthly trades conditional on trading in that month. The sample period ranges from May 2018 to April 2019. The last column shows the average maturity (in days) of the contracts that are traded by the firms of different types. Only trades between dealers and other firms have been considered when calculating the statistics. Numbers are rounded to the nearest integer or to the nearest hundred thousand.

trader type	# traders	avg. notional	CPs/month	trades/month	avg. maturity
CENTRAL BANK	45	35,013,323	5	41	24
EMPTY	14,215	13,395,139	3	144	42
FUND	11,055	10,155,767	4	1,190	38
GOVERNMENT	94	41,151,354	13	1,118	43
INSURANCE & PENSION	524	66,889,994	10	406	36
NON-FINANCIAL	6,739	9,458,687	7	1,349	59

Virtually no trades in the FX forward dealer-to-customer (D2C) market are cleared through CCPs. Of almost 3 million trades, we only have less than 500 trades involving CCPs. Table 1 considers all customers that are not CCPs and shows characteristics of D2C trades in our sample for each type of trader. The average notional of trades by insurers and pension funds is largest. The mean notional volumes of trades by funds and nonfinancial firms are considerably smaller (EUR 10.2 million and EUR 9.5 million, respectively). As traders, funds have on average 4 counterparties in a month conditional on trading in the first place. Firms classified as EMPTY, have even fewer counterparties with on average 3 trading partners per month conditional on trading in that month. Nonfinancial firms have more counterparties per month, with an average number of monthly trading partners of almost 7. Governments have the highest number of monthly counterparties. Similar comments apply to the average number of monthly trades conditional on trading in the first place. Strikingly, despite having the largest average notional volume per trade, insurance firms and pension funds are associated with a small number of monthly trades compared to funds or non-financial firms. One can also see in the last column of Table 1 that there is some disperion in the types of contracts traded across the different types of traders. While central banks rather trade contracts with a short maturity (24 days), funds, governments and nonfinancial firms trade contracts with longer maturities (38, 43 and 59 days, respectively). Figure A.1 shows the distribution of maturities across all traded contracts.

Firms that act as dealers in the FX forward market are labelled either as banks or as G16 dealers. Table 2 shows the same statistics considered in Table 1 for the two types of dealers. The average notional of trades involving G16 dealers is roughly EUR 15 million which is considerably larger than the average notional of roughly EUR 7 million of trades between other banks and their clients. Consistent with the core-periphery structure described in Abad et al. (2016) and analogous findings for other OTC markets, G16 dealers have a lot more monthly counterparties and on average much more monthly trades than other banks. The average maturities of the contracts traded by the two different types of intermediaries are relatively similar.

Table 2: Dealer characteristics in the D2C market. This table shows the number of firms for each type in the sample, shows how large the notional volume in EUR per trade involving these firms is and looks at the dealer characteristics for trades involving these dealers. The dealer characteristics are the number of counterparties and the number of trades in the sample period from May 2018 to April 2019. The last column shows the average maturity of the contracts that are traded by the dealers of different types in days. Only trades between dealers and other firms have been considered when calculating the statistics. Numbers are rounded to the nearest integer or to the nearest hundred thousand.

trader type	# dealers	notional/trade	CPs	trades	avg. maturity
BANK	201	7,232,527	821	42,824	42
G16	16	14,672,750	4,553	208,195	44

This paper mostly focuses on the D2C market for two reasons. First, assuming D2C trades are client-initiated allows us to sign these trades. Since the EMIR dataset does not indicate which counterparty initiates the trade, it is harder to sign client-to-client (C2C) or dealer-to-dealer (D2D) trades. Second, the C2C market is not very active. While a limited number of firms generated a high number of trades in the sample period, notional volumes tend to be small. Thus, compared to the notional volume in the D2D and D2C markets, the C2C market is small. Table A.1 in Appendix A shows how many trades were executed between the different groups of traders in the C2C market and Table A.2 in Appendix A shows the corresponding average notional volumes.

The D2D market, on the other hand, is very large. Table A.3 in Appendix A shows how many trades were executed between the different types of dealers in the D2D market and how much notional was exchanged on average.

Table 3 breaks down the trading done by different types of counterparties with the two sets of dealers. The most striking feature is the predominance of trading with G16 dealers as opposed to with smaller banks. In terms of notional volume G16 dealers execute around 80% of the volume. In terms of the number of trades we see more diversity. For instance, government entities execute 94% of their notional trades with G16 dealers, while that share falls to 62% for nonfinancial firms.

Table 3: Who trades with whom in the D2C market? This table shows how much notional volume in EUR a trader of each type trades on average with G16 dealers and other banks, respectively, as well as how many transactions happen on average between a trader of a given type and G16 dealers of other banks, respectively. Notional values are rounded to the nearest million. Numbers of trades are rounded to the nearest integer.

	notional volume		# trades	
	total	% traded with G16	total	% traded with G16
CENTRAL BANK	3,001mn	74%	87	85%
EMPTY	436mn	84%	33	77%
FUND	1,142mn	87%	112	79%
GOVERNMENT	8,900mn	85%	221	94%
INSURANCE & PENSION	8,302mn	89%	125	83%
NON-FINANCIAL	782mn	82%	80	62%

2.2 Benchmark rates

We use data from the Thomson Reuters Tick History database in order to calculate benchmark forward rates. We follow the same procedure to compute a benchmark for the spot rate and the forward adjustment separately. Specifically:

1. For each second, the best bid and ask prices among all dealers are determined
2. In case there are no observations in a second, the benchmark price from the previous second is used. However, a given price can only be carried forward for 30 consecutive seconds.
3. In each second, the benchmark is the average of the best bid and ask.

The final benchmark forward rate for a given tenor in a given second is the sum of the benchmarks for spot rate and forward adjustment for a specific tenor. The tenor can be overnight, 1 week, 2 weeks, 3 weeks, 1 month, 2 months, 3 months, 6 months, 9 months or 1 year. In order to obtain the benchmark rates for the forward contracts in the EMIR dataset, we use linear interpolation between the two nearest-maturity benchmark rates. We use the same procedure to calculate the benchmark rates 1, 5 or 30 minutes after each transaction.

2.3 Volatility

As exchange-rate volatility measure we use an exponentially weighted moving average of squared returns of the one-week forward exchange rate from one second to the next, i.e.

$$volatility_t^2 = 0.001 \times ret_t^2 + 0.999 \times volatility_{t-1}^2,$$

where ret_t refers to the one-second return (between t and $t+1$) on the one-week forward exchange rate. This measure captures short-lived fluctuations in volatility within a day. Using a different maturity forward rate (instead of one-week) will not significantly affect this measure as short-run fluctuations in the forward exchange rates are mostly driven by the spot exchange rate.

2.4 Price impact

The 1-minute price impact is defined as the 1-minute change in the benchmark rate times the direction of the trade (+1 if it is a client-buy and −1 if it is a sell). Analogously, we calculate the x-day permanent price impact as the difference between the benchmark rate at the time of a transaction and the end-of-day benchmark rate x-days later times the direction of the trade.[4]

Table A.4 shows the average price impact of trades by the different groups of market participants along with the corresponding standard deviations for the 1-minute and 1-day horizons. We see that 1-minute price-impact tends to be positive and highly statistically significant for all trader groups except for central banks. A positive price impact implies that clients buy (sell) on average when benchmark rates increase (decrease) subsequent to their trade. Such situation should arise if clients are on average better informed than dealers about future exchange rate movements, that is if dealers face adverse selection. This is the mechanism presented in the traditional microstructure literature (e.g., Kyle (1985) and Glosten and Milgrom (1985)). Of course, it is perhaps less intuitive to think that clients have private information about future exchange rate movements, which we typically think of as reflecting macro-economic risks. However, as we show in the model section, clients' information may pertain to their individual order flow, which may be correlated to total order flow, which in the short run may affect the change in exchange rates (see also Evans and Lyons (2002, 2005)). Alternatively, some investors may also be better at interpreting public news, say about macroeconomic fundamentals, and thus effectively also have private information about systematic sources of risk (e.g., Kim and Verrecchia (1991)). Interestingly, we see that, at a 1-day horizon, price impact remains positive and highly statistically significant on average for all traders except for Insurance & Pension trader types who display negative price impact, which effectively implies that at the longer horizon their trades on average tend to loose money, as one might explain if their trading were motivated by hedging motives for example.

Table A.5 shows the price impact of the C2D trades aggregated at the dealer level for G16 and Banks separately. We find strong evidence that dealers face adverse selection both at the 1-minute and 1-day horizon, as price impact is positive and highly statistically significant in all cases.

Of course, if dealers expect to incur a price impact cost on their client trades, it would be natural for them to charge an ex-ante premium, a 'markup,' to account for this risk. We next explain how we compute markups on C2D trades.

2.5 Markups

We define a trade's markup as the difference between the transaction rate and the benchmark rate times the direction of the trade. Table A.6 shows the average markups for the different trader types and their

[4]For that calculation we hold interpolation weights fixed and for each tenor, use the last quoted price before 8pm UCT. We further ignore week-ends that is treat the data as if Mondays follow Fridays.

respective standard deviations. We see that markups tend to be positive and statistically significant for all trader types. There seems to be an interesting positive relation between price impacts and markups, in that trader types that have higher price impact typically tend to be charged higher markups. For example, central banks face the smallest markup and hedge funds the highest. However, the relation is not monotone as non-financial traders face high markups (even higher than funds on average) and there seems to be substantial cross-sectional variation in markups. Table A.7 shows the markups aggregated at the dealer type. G16 dealers charge on average significantly smaller markups then non G16 banks, but there is a lot of variation across dealers. Figure A.2 in Appendix A shows the time series of daily average markups across all trades. The distribution of markups does not seem to exhibit any trends. In the next sections, we take a closer look at the determinants of price impact and markups, and at the relation between both.

3 Informed Clients

As shown in Tables A.4 and A.5, there is evidence that some groups of traders have significant positive price impact both at the 1-day and 1-minute horizons, which suggests that some traders have better information about future exchange rate changes. However, there is also substantial cross-sectional variation in measured price impact across traders and over time. In this section, we investigate if there are persistent differences in price impact across traders. That is, if we can find evidence that some (groups of) traders are consistently better than others at predicting future exchange rates, in the sense that they earn consistently significantly higher trading profits. To be more specific, suppose that the price impact of trader i at time t is given by

$$PI_{it} = \mu_i + \varepsilon_{it},$$

where ε_{it} is iid with finite variance and zero expectation and $\mu_i \in \mathbb{R}$. We would like to test if there is dispersion in μ_i across traders and specifically, whether some (groups of) traders have significantly higher $\mu_i > \mu_j$, say. The more transactions we observe for a given trader, the better our estimate of μ_i. Analogously, forming groups of traders gives us a relatively precise estimate of a group's average μ_i. In order to still have a sufficient dispersion in those averages, our number of groups cannot be too low. We choose to form 30 groups to obtain a good trade off between minimizing the error variance while retaining enough dispersion in the groups' price impacts. Lastly, we would like to form different groups of traders according to characteristics that are correlated with the μ_i, but not with the error ε_{it}. This rules out sorting traders based on their realized price impact, since this measure is correlated with the error. Instead, we sort traders based on their number of trades, since this is likely uncorrelated with the error, but potentially correlated with skill μ_i.

To be specific, we proceed as follows. Considering only traders that traded in both halves of our sample period, we sort the traders based on the number of trades done in the first half of the sample. Then we keep adding the traders to a group until the total number of trades in that group exceeds 1/30 times the total number of trades in the first half of the sample period. We then start adding the next traders to a group until the total number of trades of group 1 and 2 exceeds 2/30 times the total number of trades in the first half of the sample period. We continue until we have sorted the traders into 30 groups. For each group and each half of our sample period, we calculate the average permanent price impact of all trades made by that group. In Panel A of Figure 1, the average 1-minute price impact in the second half of the sample (PI2) of each group is plotted against the corresponding average 1-minute price impact in the first half of the sample (PI1). In Panel B, the same is done using the 1-day price impact. Both Panels of Figure 1 suggest that price impact is persistent.

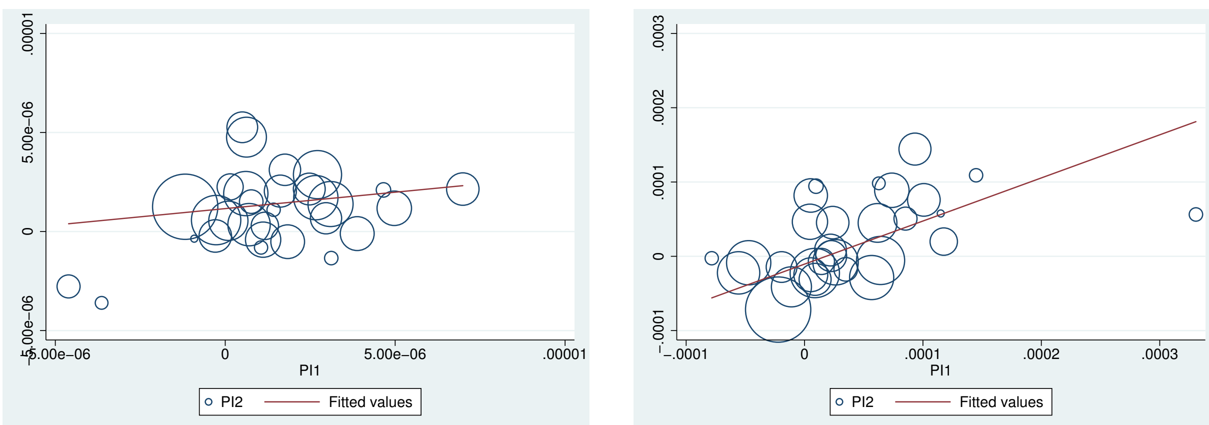

Panel A (1 minute). Panel B (1 day).

Figure 1: Persistence of price impact. In Panel A, a group's average 1-minute price impact in the second half of the sample is plotted against its 1-minute price impact in the first half of the sample. In Panel B, a group's average 1-day price impact in the second half of the sample is plotted against its 1-day price impact in the first half of the sample. Larger circles correspond to groups in which traders generate a higher notional volume (EUR) per trade in the first half of the sample. The red lines show the fitted values of linear regressions.

Different traders may have different investment horizons. In particular, there may be a set of traders whose goal it is to trade intraday on very short-lived signals. Even though these market participants might trade very profitably, their 1-day price impact may not be very persistent, since it is not their objective to trade on long-term price changes. On the other hand, we would expect strong persistence in 1-minute price impact if there are informed traders in this set. Figure A.3 in Appendix A shows the distribution of the traders' average numbers of trades per day (conditional on trading on a day). Roughly 1% of traders do more than 10 trades per day if they trade at all. We classify these traders as HFT. To investigate whether there are differences between the price impact of the HFT traders and the lower-frequency traders, we compare the persistence of the 1-minute price impact of HFT traders in Panel A of Figure 2 with that of non-HFT traders in Panel B. We find strong evidence - indeed stronger compared to Panel A of Figure 1 - of persistence in 1-day price impact for HFT traders (Panel A), and hardly any evidence of persistence for non-HFT traders in Panel B.

One can see in both panels that a high price impact is not related to high notional volume per trade, as it might be the case if the price impact were inventory-driven. Such an inventory-based price impact may arise as follows. If a dealer takes a large customer order, this is a private transaction between a dealer and its client. Other dealers will not change their quotes at the very time of the transaction. But shortly after the trade, the dealer who took the customer order may try to offset the inventory shock in the interdealer market, leading other dealers to change their quotes as well.

It seems that traders with high price impact typically have a lower average notional per trade. Similarly, Figure A.5 in the appendix shows that groups with high price impact are also not groups with high total notional traded. Thus, price impact seems more likely to be information-based.[5]

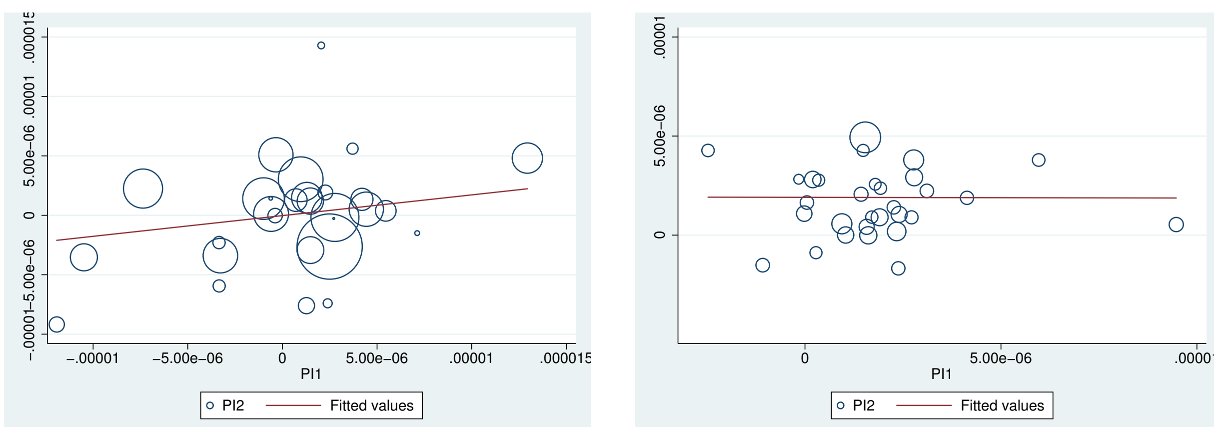

Panel A (HFT). Panel B (non-HFT).

Figure 2: Persistence of 1-minute price impact for HFT vs. non-HFT market participants. In both panels, a group's average 1-minute price impact in the second half of the sample is plotted against its 1-minute price impact in the first half of the sample. In Panel A, only HFT market participants are considered and in Panel B, only non-HFT market participants are considered. Larger circles correspond to groups in which traders generate a higher notional volume (EUR) per trade in the first half of the sample. The red lines show the fitted values of linear regressions.

[5]However, this information may very well be information about aggregate inventory or order flow.

If instead we focus on the 1-day price impact, the opposite picture emerges. Figure 3 shows that the 1-day price impact is more persistent for non-HFT traders (Panel B) than for HFT traders (Panel A). Neither in Figure 2 nor in Figure 3, is it the case that high price impact is associated to high notional volume per trade. Figures A.6 and A.7 in the appendix show that, for the same groups of traders, high price impact is not related to high total notional volume traded. Again, this is consistent with information-driven price impact, but inconsistent with inventory-based explanation of persistently positive price impact.

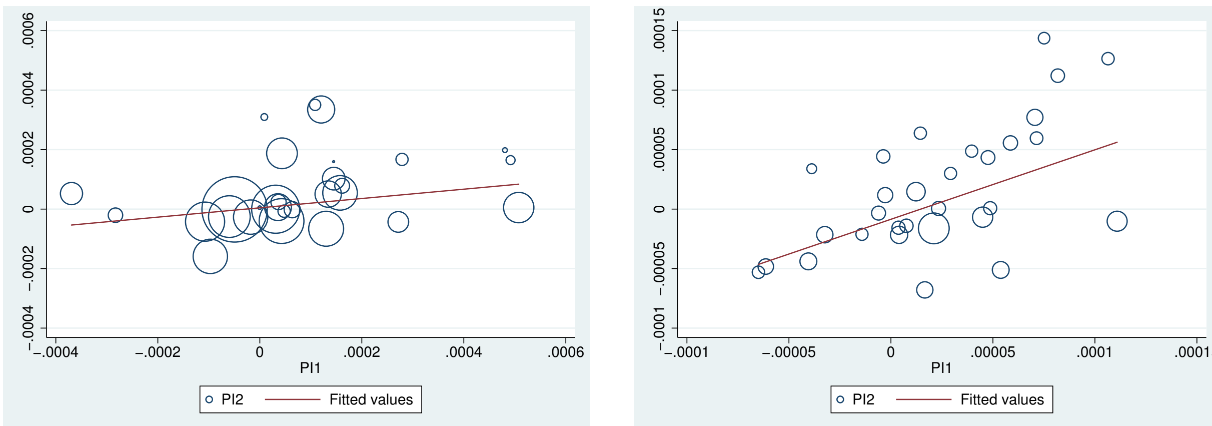

Panel A (HFT). Panel B (non-HFT).

Figure 3: Persistence of 1-day price impact for HFT vs. non-HFT market participants. In both panels, a group's average 1-day price impact in the second half of the sample is plotted against its 1-day price impact in the first half of the sample. In Panel A, only HFT market participants are considered and in B, only non-HFT market participants are considered. Larger circles correspond to groups in which traders generate a higher notional volume (EUR) per trade in the first half of the sample. The red lines show the fitted values of linear regressions.

In order to formally assess the persistence of the price impact in Figures 1 to 3, we regress a group's average price impact in the second half of the sample on its average price impact in the first half of the sample. The results for the various groups of traders and different horizons are shown in Table 4. One can see that estimates for the coefficient in front of the price impact in the first half of the sample are statistically significant except for the 1-minute price impact of non-HFT traders. Moreover the R^2 statistics from the regressions are generally large, especially for non-HFT traders' 1-day price impact and for HFT traders' 1-minute price impact.

Table 4: Persistence of price impact for different traders and horizons. This table regression coefficients and robust standard errors for the regression

$$PI2 = \beta_0 + \beta_1 PI1 + \varepsilon,$$

where $PI2$ is the price impact in the first half of the sample, $PI2$ is the price impact in the second half of the sample and ε is an error term. We use the average price impact generated by the groups of traders shown in Figures 1 to 3. Columns 1 and 2 refer to the groups from Panel A and B, respectively, of Figure 1. Columns 3 and 4 refer to the groups from Panel B of Figures 2 and 3, respectively. Columns 5 and 6 refer to the groups from Panel A of Figures 2 and 3, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	all	all	non-HFT	non-HFT	HFT	HFT
	1 min	1 day	1 min	1 day	1 min	1 day
PI1	0.33**	0.34**	-0.01	0.72***	0.37***	0.20**
	(0.14)	(0.14)	(0.16)	(0.21)	(0.13)	(0.09)
Constant	0.00	0.00	0.00***	-0.00	-0.00	0.00*
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
N	30	30	30	30	29	29
r2	0.17	0.24	0.00	0.38	0.15	0.10

$p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In order to characterize informed and uninformed traders, we look at the three groups with the highest price impact in the second half of the sample in Panel B of Figure 1 and call the traders in those groups 'high-PI traders.' Analogously, we call all traders in the groups from Panel B in Figure 1 with negative price impact in both halves of the sample 'neg-PI traders.'

Table 5 shows characteristics of neg-PI or high-PI traders depending on whether they belong to the HFT group or not. One can see that for both non-HFT and HFT market participants, informedness is negatively related to trading volume and number of counterparties. On the other hand, the relationship between the number of monthly trades is nonmonotone. For the less active non-HFT market participants, more monthly trades are associated with being uninformed. On the other hand, informed HFT market participants trade more often than their less informed neg-PI counterparts. For both non-HFT and HFT market participants, trading longer maturities is associated with being in the high-PI group, i.e. higher informedness, which may be consistent with these traders seeking the largest exposure.

Table 5: Trader characteristics. The following table shows properties of neg-PI and high-PI traders depending on whether they belong to the HFT group or not. Notional values are rounded to the nearest hundred thousand. The numbers of counterparties and average monthly trades are rounded to one decimal and the average maturity is rounded to the nearest integer.

	non-HFT		HFT	
	neg-PI	high-PI	neg-PI	high-PI
notional/trade (EUR)	21.4mn	9.6mn	11.9mn	1.4mn
counterparties	4.1	1.2	6.1	3.8
avg. monthly trades	102.9	3.0	706.0	1221.0
average maturity (days)	34	61	62	69

Looking only at high-PI and neg-PI traders, Table 6 examines which trader characteristics are associated with being informed. Table 6 shows the result of a linear probability regression model to explain a dummy variable that is one if the trader is a high-PI trader and zero otherwise. A trader's number of counterparties, traded notional and the number of monthly trades are negatively related to being a high-PI trader, as one can observe in columns 1 to 3. Also, HFT market participants are more likely to have a lower 1-day price impact, as shown in column 4. Notional volume is not significant anymore in column 5, when controlling for the number of trades and the number of counterparties. One can also see that the number of trades has different implications for the probability of being informed depending on whether the trader belongs to the HFT group or not. The number of counterparties is still negatively related to being informed even when controlling for other factors. This result seems to run against the findings of Kondor and Pintér (2019) that informed traders have more counterparties, but is consistent with information leakage examined empirically in Hendershott and Madhavan (2015) and Hagströmer and Menkveld (2019) and modeled theoretically in Liu et al. (2018). The cost of information leakage is higher for informed traders, which may be the reason why they contact fewer dealers.

Table 6: Probability of being an informed trader. This table shows coefficient estimates B and robust standard errors for the regression

$$informed = BX + \varepsilon,$$

where $informed$ is a dummy variable that is equal to one if a trader belongs to the high-1-day-PI group and equal to zero otherwise. The vector X includes different trader characteristics specified in the table and ε is an error term. The sample includes all high-Pi and neg-PI traders from the groups in Figure 1. The averages of the numbers of trades have been divided by 10^3 and averages of the notionals per trade have been divided by 10^{11}.

	(1)	(2)	(3)	(4)	(5)
avg. monthly counterparties (CPs)	-84.70***				-12.07***
	(12.17)				(3.22)
avg. monthly trades		-0.19**			-7.87***
		(0.09)			(0.23)
avg. monthly trades × HFT dummy					7.91***
					(0.23)
avg. notional in EUR			-23.31***		4.82
			(7.95)		(5.99)
HFT dummy				-0.52***	-0.57***
				(0.08)	(0.08)
Constant	1.07***	0.96***	0.96***	0.96***	1.03***
	(0.02)	(0.00)	(0.00)	(0.00)	(0.00)
N	9628	9628	9628	9628	9628
r2	0.26	0.03	0.00	0.03	0.80

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

To see how robust our findings about the persistence of the price impact of different traders, we break the sample into four different subperiods, and check whether traders who had a high price impact in the first period also have a high price impact in the following periods. To this end, we sort all traders who were active in all four quarters of our sample into two groups based on their number of trades in the first quarter and then compute their average price impact in each quarter of the sample. Figure 4 shows that the traders who had a high price impact in the first quarter generally also have a higher price impact in the subsequent quarters. Moreover, if one restricts attention to non-HFT market participants, traders who have a higher price impact in the first quarter also have a higher price impact in the last quarter.

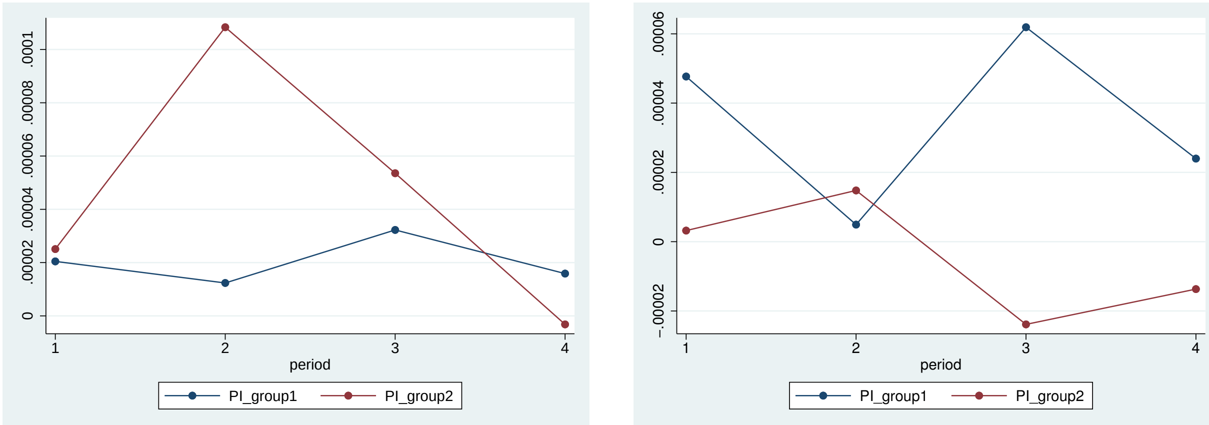

Panel A (all traders).

Panel B (non-HFT).

Figure 4: Persistence across multiple subperiods. We sort all traders who were active in all four quarters of our sample into two groups based on their number of trades in the first quarter and then plot their average 1-day price impact in each quarter of the sample. In Panel A, transactions by all traders were considered and in Panel B, only trades by non-HFT market participants were considered.

4 Informed Dealers

Evans and Lyons (2002) find that aggregate orderflow in the interdealer market predicts future price changes in the DM/USD spot market using four months of data from an interdealer trading system in 1996. Most of the variation in forward rates comes from variation in the spot rate. Moreover, order flow in the interdealer market is (in most models, like the one presented in this paper) generated by order flow from customers. Thus, one may hypothesize that customers' order flow predicts future price changes in the EUR/USD forward market. The model presented in Section 6 presents a mechanism that would generate such predictability and also suggests that the customer order flow of different dealers has different predictive power for future price changes. Similar to Evans and Lyons (2002), we look at the order imbalance, i.e. the difference between daily buy and sell orders received by a dealer. We define the volume imbalance as the difference between notional volume bought by customers and notional volume sold by customers on a given day. Table 7 describes order imbalance and volume imbalance for the different dealer types. One can see that both kinds of dealers have relatively small average order imbalances and small average volume imbalances. However, G16 dealers have much more volatile order imbalances and volume imbalances as indicated by their larger mean absolute order and volume imbalances.

Table 7: Order imbalance (OI). This table shows the average order imbalance, the average of the absolute values of order imbalances for each type of dealer as well as the respective values for the volume imbalance.

trader type	mean OI	mean absolute OI	mean volume imbalance	mean absolute volume imbalance
BANK	-3	8	-2.6mn	58.1mn
G16	-2	66	3.6mn	809.0mn

In order to measure how well each dealer can forecast future price changes using their customers' order flow, we regress changes in the benchmark price for the one-week forward exchange rate between the end of day t and the end of day $t+1$ on the sum of the order imbalances on the last five days on which the dealer has traded. The higher the resulting R^2 statistic, the more informed is the dealer. As we show in the model in section 6.3 the R^2 statistic should also be a determinant of a dealer's markups as it is related to a dealer's expectation of the price impact it will incur.

Table 8 shows summary statistics of the R^2 we obtain for every dealer in the sample, provided that enough data is available to perform the regression described above. We also transform the R^2 into an annual Sharpe ratio using the method in Table 1 in Cochrane (1999), assuming the true mean of price rate movements is zero and 250 trading days in a year.[6] One can see in Table 8 that even though the R^2 statistics obtained in *daily* forecasts seem small, annualized Sharpe ratios are actually substantial.

[6] Cochrane (1999) derives the formula

$$SR^{informed,\,annual} = \frac{\sqrt{(SR^{uninformed,daily})^2 + R^2}}{\sqrt{1-R^2}}\sqrt{days/year},$$

where $SR^{informed\,annual}$ is the annualized Sharpe ratio, a dealer that has a given R^2 when forecasing daily returns can achieve, given that an investor not forecasting returns can achieve a daily Sharpe ratio of $SR^{uninformed\,daily}$. Under the assumptions in the text, this formula becomes

$$SR^{informed,\,annual} = \frac{\sqrt{R^2}}{\sqrt{1-R^2}}\sqrt{250}.$$

Table 8: Dealer informedness and Sharpe ratios. This table shows summary statistics for the R^2 statistic from the regression of price changes on the dealers order imbalances described in the text, i.e.

$$\frac{rate_{t+1} - rate_t}{rate_t} = \beta_{0,i} + \beta_{1,i} sum\_OI_i + \varepsilon_{it},,$$

where ε_{it} is an error term, sum_OI_i refers to the sum of the order imbalances of dealer i for its last 5 trading days and $rate_t$ refers to the one-week forward exchange rate at day t.

Variable	Obs	Mean	Std. Dev.	Min	Max
R^2	142	0.022383	0.047305	.000002	0.27981
Sharpe ratio	142	1.72584	1.868606	0.020974	9.855478

In order to study informed and uninformed dealers, we look at trade-by-trade data and assign trades to two quantiles according to the informedness of the dealer trading. We create a dummy variable that is equal to 1, if the dealer informedness is above the median informedness across all trades. For all other trades, the dummy variable is equal to zero. We call dealers for which this dummy variable is one "informed."

Some of the dealers in our sample trade very infrequently. Especially those dealers with extreme Sharpe ratios only have a very limited number of days on which they trade in our sample. Figure A.8 in Appendix A shows the distributions of the R^2 across dealers. Most dealers have an R^2 of less than 0.025. In order to avoid focussing on outliers when studying the characteristics of informed dealers, we focus on dealers with an R^2 of less than 2% and a notional trading volume of at least 0.5% of the entire market. As an additional robustness check we also look separately at dealers with a notional trading volume greater than 2.5% of the entire market.

In Table 9 one can see the characteristics of the high versus low R^2 dealers, depending on the fraction of total D2C volume in EUR they are responsible for. We see that more informed dealers tend to have smaller notionals per trade and have fewer counterparties in D2C and D2D markets.

Table 9: Dealer characteristics. The following table shows properties informed and uninformed dealers depending on how much notional volume they trade. Percentages are rounded to one decimal, numbers of counterparties are rounded to the nearest integer and notional volumes are rounded to the nearest hundred thousand.

	2.5% > volume > 0.5%		volume > 2.5%	
	uninformed	informed	uninformed	informed
% G16	66.6%	16.7%	100%	100%
dealer's avg. notional/trade	12.1mn	9.2mn	17.5mn	16mn
D2C counterparties	1399	1288	5106	3947
D2D counterparties	123	102	266	220
% of total notional D2C volume (in EUR)	5.2%	8.7%	50%	30.3%

Table 10 shows the result of a linear probability regression model , which predicts the informdness dummy based on various dealer characteristics. One can see that both traded notional and the number of trades are highly significant and negatively correlated with being informed. This is true for both subsets of dealers. In the regressions shown in Table A.11 in Appendix A we use the dealer's R^2 as the left-hand side variable instead of the informedness dummy and get similar results.

Table 10: Probability of being an informed dealer. We run the regression

$$info\ dummy = BX + \varepsilon,$$

where $info\ dummy$ is the dummy variable described in the text which measures the informedness of a dealer. This table shows the coefficient estimates B for various explanatory variables X as well as robust standard errors. In columns 1 to 3 we focus on dealers executing more than 0.5% of the notional volume (EUR) of the entire D2C market. In columns 4-6 we focus on dealers executing more than 2.5% of the notional volume of the entire D2C market. The number of trades has been divided by 10^7 and the notional traded has been divided by 10^{10}.

	> 0.5% notional volume			> 2.5% notional volume		
	(1)	(2)	(3)	(4)	(5)	(6)
avg. notional traded	-202.75		-333.44*	-149.26		-807.10***
	(151.30)		(189.49)	(258.89)		(92.44)
# trades		-24.36**	-41.15***		-36.34***	-62.32***
		(8.60)	(12.92)		(10.81)	(7.64)
D2D counterparties			0.00			
			(0.00)			
D2C counterparties			0.00			
			(0.00)			
Constant	0.88***	0.90***	1.23***	0.80	1.18***	2.99***
	(0.24)	(0.18)	(0.29)	(0.50)	(0.28)	(0.28)
N	20	20	20	11	11	11
r2	0.06	0.20	0.30	0.02	0.43	0.87

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure A.4 in Appendix A shows that traders with higher price impact tend to pay higher markups. This suggests that dealers may price discriminate based on client identities or characteristics. Since there is evidence that clients' price impact is persistent, it would be natural to think that dealers set markups based on past client price impact. At the same time, since we have shown that dealers are differentially informed about future price changes based on their clients' aggregated order imbalance, it is also natural to think that dealers will use that information to set markups. To better understand the determinants of markups, we now perform a panel regression, where we explain client markups with various client, dealer, and trade characteristics.

To eliminate data errors, we exclude all trades with markups below -2% and above 3%. This effectively puts a 2% band around the traders' average markups shown in Figure 5 in Hau et al. (2019). In column 1 of Table 11, the coefficient for realized values of 1-day price impact is significant at the 5% level. Moreover, we find a significant interaction term with the informedness dummy, which suggests that informed dealers' markups react more strongly to the realized 1-day impact. This is consistent with the implications of our model in section 6.3, where dealers that have more informed order-flow are better at predicting future price changes and thus set markups that are more in line with future price changes.

Realized values of price impact measures may be significant in the regression from Table 11, because they are correlated with other trader characteristics. In order to control for those, we add fixed effects for each dealer-trader pair along with other time-varying control variables. The results of this regression are shown in column 2 of Table 11. The coefficients on the realized 1-day price impact and the interaction term in the first row are very similar to the estimates in column 1 and still significant.

In order to asses how much the results in columns 1 and 2 of Table 11 are driven by the dealers' connections in the D2D market, we replace the informed-dealer dummy by a connected-dealer dummy in column 3. To this end, we assign trades into two quantiles according to the number of the dealer's counterparties in the D2D market. The connected-dealer dummy is a dummy variable that is equal to 1, if the dealer's number of D2D counterparties is above the median across all trades. The results of this regression are shown in column 3 of Table 11. We find that more connected dealers respond less to changes in future prices, as measured by realized values of 1-day price impact. However, the coefficient in front of the interaction term is not even significant at the 10% level, whereas the coefficient in front of the interaction term in column 2 of Table 11 is significant at the 5% level. It thus seems that the informed dealer dummy better captures differences in the dealers' sensitivity with respect to future price changes.

In column 4 of Table 11, we study what may drive potential dealer and trader fixed effects on markups. One can see that traders that have on average a higher 1-day price impact also have to pay higher markups. The more counterparties a trader has, the lower the markups possibly due to increased bargaining power.

Table 11: Markups and price impact. We run the regression

$$markup_{it} = BX + \varepsilon,$$

where ε_{it} is an error term, X are explanatory variables specified in the table and $markup_{it}$ is the markup that trader i has to pay at time t. We report coefficients B and standard errors that are clustered at the dealer level. We excluded all trades with markups below -2% and above 3%. Order imbalance and its standard deviation have been divided by 10^6. Lastly, the logs of counterparties and trades have been divided by 1000.

	(1)	(2)	(3)	(4)
realized 1-day impact × informedness dummy info	0.0112***	0.0099**		0.0104**
	(0.0043)	(0.0043)		(0.0044)
realized 1-day impact × connectedness dummy			-0.0068	
			(0.0042)	
realized 1-day impact	0.0107***	0.0091***	0.0171***	0.0091***
	(0.0012)	(0.0015)	(0.0031)	(0.0015)
realized 1-minute impact		0.0949***	0.0892***	0.0837**
		(0.0330)	(0.0337)	(0.0338)
market conditions:				
volatility		0.7146*	0.7286*	1.0947*
		(0.3640)	(0.3804)	(0.6527)
Smart average 1-day impact group		0.0013	0.0019	0.0035
		(0.0061)	(0.0064)	(0.0073)
time-varying trader characteristics:				
log(traders' monthly counterparties)		-0.0139	-0.0145	-0.0901***
		(0.0212)	(0.0213)	(0.0215)
log(traders' monthly trades)		0.0269***	0.0276***	-0.0319
		(0.0086)	(0.0086)	(0.0238)
time-varying dealer characteristics:				
dealer's signed OI		0.0350	0.0343	0.0326
		(0.0323)	(0.0323)	(0.0339)

(To be continued)

Table 11-Continued.

	(1)	(2)	(3)	(4)
fixed trader characteristics:				
trader's average 1-day impact				0.1098***
				(0.0173)
high-PI dummy				0.0002***
				(0.0001)
neg-PI dummy				-0.0001*
				(0.0000)
HFT				-0.0000
				(0.0001)
fixed dealer characteristics:				
informedness dummy	0.0000			-0.0000
	(0.0001)			(0.0001)
connectedness dummy				0.0000
				(0.0001)
standard deviation of dealer's OI				-0.0983*
				(0.0499)
Constant	0.0002***	0.0000	0.0000	0.0004**
	(0.0000)	(0.0001)	(0.0001)	(0.0002)
dealer-client fixed effects	no	yes	yes	no
N	2770512	2742738	2684941	2684941

$p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

To examine how trades by very active traders are different from those trades made by less active market participants, we focus only on HFT market participants in column 1 of Table 12. We can see that results related to 1-day price impact are less strong than for the whole sample, but the coefficient on 1-minute price impact is larger. On the other hand, the results associated with 1-day price impact are stronger when focussing on non-HFT market participants. The results are not much affected if we exclude neg-PI traders, or if we distinguish between low- or high-volatility periods. However, comparing columns 5 and 6 in Table 12, we see that traders with a high average price impact have to pay even higher markups in high-volatility periods.

Table 12: Markups and price impact in different subsamples We run the regression

$$markup_{it} = BX + \varepsilon,$$

where ε_{it} is an error term, X are explanatory variables specified in the table and $markup_{it}$ is the markup that trader i has to pay at time t. We report coefficients B and standard errors that are clustered at the dealer level. In column 1, we focus on HFT market participants, and in column 2 we focus on non-HFT market participants. In column 3 we exclude all traders labelled as neg-PI traders. In column 4 we focus only on those traders. In column 5 we focus on the trades for which volatility is below the median across all trades and in column 6 we focus on trades for which volatility is above the median across all trades. We excluded all trades with markups below -2% and above 3%. Order imbalance and its standard deviation have been divided by 10^6.

	(1)	(2)	(3)	(4)	(5)	(6)
	HFT	non-HFT	no neg-PI	neg-PI	low vol	high vol
real. 1-day impact × info dummy	0.006	0.012**	0.009**	0.020***	0.011*	0.009***
	(0.004)	(0.005)	(0.004)	(0.007)	(0.006)	(0.003)
realized 1-day impact	0.010***	0.009***	0.009***	0.007***	0.009***	0.009***
	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
realized 1-minute impact	0.139***	0.054*	0.078***	0.116	0.089***	0.081
	(0.053)	(0.029)	(0.027)	(0.088)	(0.030)	(0.055)
market conditions:						
volatility	2.399**	0.348	1.197*	0.284	2.551***	-0.870
	(1.003)	(0.501)	(0.708)	(0.512)	(0.674)	(1.344)
Smart average 1-day impact	0.005	0.002	0.005	-0.004	-0.004	0.011**
	(0.010)	(0.008)	(0.007)	(0.014)	(0.012)	(0.005)
varying trader characteristics:						
log(traders' monthly counterparties)	-0.033**	-0.131***	-0.101***	-0.042**	-0.088***	-0.095***
	(0.016)	(0.035)	(0.023)	(0.019)	(0.022)	(0.023)
log(traders' monthly trades)	0.008	-0.068***	-0.030	-0.018	-0.049**	-0.016
	(0.023)	(0.022)	(0.025)	(0.017)	(0.021)	(0.025)
varying dealer characteristics:						
dealer's signed OI	0.050	0.019	0.037	0.006	0.025	0.040
	(0.043)	(0.026)	(0.036)	(0.024)	(0.025)	(0.043)

(To be continued)

Table 12-Continued.

	(1)	(2)	(3)	(4)	(5)	(6)
	HFT	non-HFT	no low-PI	low-PI	low vol	high vol
fixed trader characteristics:						
trader's average 1-day impact	0.043	0.117***	0.112***	-0.059	0.099***	0.121***
	(0.071)	(0.018)	(0.017)	(0.084)	(0.021)	(0.027)
high-PI dummy	0.000**	0.000**	0.000***	0.000	0.000**	0.000***
	(0.000)	(0.000)	(0.000)	(.)	(0.000)	(0.000)
neg-PI dummy	0.000	0.000	0.000	0.000	-0.000	-0.000**
	(0.000)	(0.000)	(.)	(.)	(0.000)	(0.000)
HFT	0.000	0.000	-0.000	-0.000	0.000	-0.000
	(.)	(.)	(0.000)	(0.000)	(0.000)	(0.000)
fixed dealer characteristics:						
informedness dummy	0.000	-0.000	-0.000	0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
connectedness dummy	0.000***	-0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
standard deviation of dealer's OI	-0.040	-0.103	-0.114**	-0.017	-0.111**	-0.084*
	(0.040)	(0.069)	(0.055)	(0.048)	(0.054)	(0.048)
Constant	-0.000	0.001***	0.000**	0.000	0.000***	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
N	905843	1779098	2274265	410676	1347896	1337045

$p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.12 in Appendix A shows similar regressions as those from Table 12, but using deale and trader fixed effects. The coefficient estimates in Table A.12 are very similar to those shown in Table 12.

Figure A.9 in Appendix A illustrates the findings from the regression with a simple plot that shows how more informed dealers charge higher markups when the price impact is higher.

Appendix C discusses the possible errors-in-variables problem that may arise because both the measured price impact and the measured markup are affected by errors in the benchmark rate. Such errors may arise because of noisy quotes in the TRTH database or imprecise time-stamps in the EMIR database, which would lead to a mechanical positive correlation between markups and realized price impact. As we discuss in the appendix, this bias is unlikely to affect our main results however.

5 Endogenous Dealer Choice

The model presented in Section 6.4 suggests that informed traders are more likely to trade with dealers who have informative order flow. In order to test this hypothesis, we study how a trader's average 1-day price impact affects the probability of trading with an informed dealer.

As in sections 3 and 4, HFT market participants arguably behave differently from non-HFT market participants. For this reason, we consider these two groups separately. In Table 13 one can see that non-HFT traders are more likely to trade with informed traders if their 1-day price impact is higher on average. Even 1-minute price impact measures positively affect the probability of trading with an informed dealer, as can be seen in column 2. These statements remain true when controlling for other trader characteristics that are associated with being informed (see Section 3), like the number of counterparties, or the number of trades. On the other hand, 1-day price impact measures do not affect dealer choice for HFT market participants, since they likely have a much shorter investment horizon. However, for HFT average 1-minute impact still positively affects the probability of trading with an informed dealer. One can also see that non-HFT clients are more likely to trade with informed dealer if volatility is high, while HFT clients are more likely to trade with uninformed dealers if volatility is high.

The results shown in Table 13 suggest that traders are more likely to trade with informed dealers if they are more informed on average (i.e. have a higher average 1-day price impact), when adverse selection is higher on average (volatility is higher), or when they are particularly well informed (the realization of the 1-day price impact turns out to be high). To avoid the endogeneity problem that we find that informed traders choose informed dealers, because we define an informed dealer based on the informativeness of her clients' order flow, we now define dealer informedness as the informativeness of the order flow from non-financial companies. This means we perform the same steps as in Section 4, including the definition of a dummy variable, but use order flow only from non-financial customers to predict future price changes. We then look at all traders except for non-financial firms and examine the probability of trading with an informed dealer. The results are shown in Table 14. We see that non-HFT traders are still more likely to trade with dealers that are informed according to this new measure.

Table 13: Probability of trading with informed dealers. We run the regression

$$informed = BX + \varepsilon,$$

where ε_{it} is an error term, X are explanatory variables specified in the table and $informed$ is the informed-dummy discussed in Section 4. We report coefficients B and standard errors that are clustered at the trader level. The dataset includes D2C trades. We excluded all trades with markups below -2% and above 3% .

	excluding HFT			HFT only		
	(1)	(2)	(3)	(4)	(5)	(6)
avg. 1-day impact	10.83***	10.81***	8.98**	-36.34	-25.24	-3.60
	(3.77)	(3.77)	(3.75)	(71.99)	(70.88)	(67.83)
avg. 1-min impact		148.05*	153.02*		3531.22	4118.06*
		(81.34)	(81.03)		(2180.98)	(2092.95)
realized 1-day impact		0.02	0.03		0.10	0.21
		(0.05)	(0.05)		(0.12)	(0.13)
realized 1-min impact		3.82***	3.88***		-0.92	-1.85
		(0.98)	(0.98)		(1.15)	(1.18)
log(monthly counterparties)			-21.45**			58.11
			(9.58)			(38.63)
log(monthly trades)			-6.75			-87.91***
			(6.63)			(29.49)
volatility			428.48***			-1563.72**
			(131.95)			(750.82)
Constant	0.54***	0.54***	0.56***	0.38***	0.38***	1.01***
	(0.01)	(0.01)	(0.01)	(0.06)	(0.06)	(0.18)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 14: Probability of trading with informed dealers without circularity. We run the regression

$$informed^* = BX + \varepsilon,$$

where ε_{it} is an error term, X are explanatory variables specified in the table and $informed^*$ is the alternative info dummy discussed this section. We report coefficients B and standard errors that are clustered at the trader level. The dataset includes D2C trades. We excluded all trades with markups below -2% and above 3% as well as trades made by non-financial firms.

	excluding HFT			HFT only		
	(1)	(2)	(3)	(4)	(5)	(6)
avg. 1-day impact	10.92**	10.77**	10.01**	221.28**	219.26**	288.62***
	(4.81)	(4.80)	(4.83)	(101.24)	(99.80)	(98.29)
avg. 1-min impact		-53.34	-37.87		-846.62	-244.94
		(98.22)	(97.69)		(3456.66)	(3409.79)
realized 1-day impact		-0.02	-0.01		-0.03	0.09
		(0.05)	(0.05)		(0.19)	(0.20)
realized 1-min impact		1.41	1.37		2.10	1.17
		(1.12)	(1.12)		(3.44)	(3.33)
log(monthly counterparties)			-23.17**			120.16***
			(10.91)			(24.41)
log(monthly trades)			4.07			-118.95***
			(7.53)			(15.32)
volatility			864.03***			-932.14*
			(150.65)			(528.57)
Constant	0.50***	0.50***	0.45***	0.30***	0.30***	1.07***
	(0.01)	(0.01)	(0.02)	(0.06)	(0.06)	(0.13)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.8 in Appendix A shows the 1-day and 1-minute price impact that trades executed by informed and uniformed dealers have. One can see that both the 1-minute price impact and the 1-day price impact are higher for trades executed by informed dealers. One obtains a slightly different result when using other proxies for dealer informedness like connections in the D2C market or connections in the D2D market. Tables A.9 and A.10 in Appendix A show that more connected dealers face a higher 1-minute price impact, but a lower 1-day price impact. In Table A.13, we run regressions similar to those in Table 13, except that we replace the informedness dummy with a connectedness dummy. In this case average price impact measures have no significant effect on dealer choice. Thus, the informedness of a dealer cannot simply be captured by the dealers D2D connections.

6 Model

The aim of this section is to lay out a mechanism how dealers learn from the orderflow of their customers and incorporate this information in their quotes. For the sake of clarity, other determinants of the dealers' quotes are not modeled explicitly. In section 6.4, the model is extended in order to account for informed trading and endogenous dealer choice.

6.1 Setup

There are $N \in \mathbb{N}$ dealers, with $N > 2$. The dealers are indexed by the set $\mathcal{D} := \{1, 2, ..., N\}$ and have two periods to trade one asset: In the first period, each dealer $i \in \mathcal{D}$ can trade with a a group of clients. The clients initiate the trade by specifying a quantity they want to trade. The dealer responds with a competitive quote at which all orders of the dealer's clients are executed. In the second period, the N dealers can trade among each other in a centralized market. Each dealer submits linear demand schedules and the price is determined by market clearing.

Let x_i denote the net amount of the asset that clients buy from dealer i.[7] These quantities are jointly normally distributed across dealers with covariance matrix Σ:

$$(x_1, ..., x_N)' \sim \mathcal{N}(0, \Sigma). \tag{1}$$

Moreover, let $p_{1,i}$ denote the price at which the transaction between dealer i and the clients happens, let τ_i denote the quantity that dealer i buys from other dealers in period 2 and let p_2 denote the price at which the dealers trade among themselves in period 2.[8]

Then, the utility of dealer $i \in \mathcal{D}$ in the end of period 2 is given by

$$U_i(x_i, p_{1,i}, p_2) := -(-x_i + \tau_i)^2 \frac{\gamma}{2} + x_i p_{1,i} - p_2 \tau_i, \tag{2}$$

where $\gamma > 0$ determines the aversion of the dealer to holding inventory. The first term in (2) represents

[7] Here, $x_i < 0$ means that the clients sell to dealer $i \in \mathcal{D}$.
[8] Again, $\tau_i > 0, i \in \mathcal{D}$ means that dealer i buys the asset.

quadratic inventory holding costs that the dealer needs to pay after period 2. The remaining two terms in (2) represent the revenue generated by trading with the clients and the money paid in the interdealer market, respectively. We assume that in period 1, dealers quote a competitive price, p_1, to all their clients such that they are indifferent between trading or not trading in the first round. In period 2, when dealers trade among themselves, dealers maximize the expectation of the expression in (2).

6.2 Equilibrium

An equilibrium in this model consists of a pricing rule for each dealer that specifies a price for a given quantity that the clients demand as well as a demand schedule for each dealer such that the interdealer market clears in period 2. Moreover, dealers behave competitively when facing clients and maximize expected utility in the interdealer market.

We conjecture that dealer $i \in \mathcal{D}$ uses a linear demand schedule $\tau_i(x_i, p_2) : \mathbb{R}^2 \to \mathbb{R}$ with

$$\tau(x_i, p_2) := ap_2 + bx_i, \tag{3}$$

where $a, b \in \mathbb{R}$. Market clearing gives the price in the interdealer market as a function of the quantities traded between dealers and their clients:

$$p_2(x_1, ..., x_N) = \frac{-b}{aN} \sum_{i=1}^{N} x_i.$$

If agent $i \in \mathcal{D}$ is strategic in the interdealer market and conjectures strategies for all the other dealers $j \neq i$ to be as in (3), then market clearing implies a residual inverse demand function such that:

$$\lambda := \frac{\partial}{\partial \tau_i} p_2 = \frac{-1}{a(N-1)}. \tag{4}$$

Maximizing the expression in (2) for a given price p_2 with respect to τ_i and calculating the corresponding the first-order condition gives the necessary condition[9]

$$\tau_i = \frac{\gamma}{\gamma + \lambda} x_i - \frac{1}{\gamma + \lambda} p_2. \tag{5}$$

Now (3), (4) and (5) imply

[9]By calculating the second derivative, one can see that this condition is also sufficient.

$$a = -\frac{1}{\gamma - \frac{1}{a(N-1)}} \tag{6}$$

and

$$b = \frac{\gamma}{\gamma - \frac{1}{a(N-1)}}. \tag{7}$$

Comparing (6) and (7) gives

$$a = -\frac{b}{\gamma}. \tag{8}$$

Using the last result and (7) gives

$$b = \frac{N-2}{N-1}. \tag{9}$$

It follows that if we define the average client demand $\bar{x} = \frac{\sum_{i=1}^{N} x_i}{N}$, then we have

$$p_2 = \gamma \bar{x} \tag{10}$$

$$\tau_i = \frac{N-2}{N-1}(x_i - \bar{x}) \tag{11}$$

We see that the dealer's trade in the inter-dealer market does not depend on her risk-aversion, but only on how different her clients' trades are from the average trade. Instead, the price level in the interdealer market is equal to the average inventory held by dealers times their risk-aversion. The price fully reveals the average inventory which is a sufficient statistic for all the dealers trades, given their own client demand.

To finish the characterization of the equilibrium, it remains to derive the prices that the dealers are offering their clients. Dealers behave competitively and achieve the same expected utility when trading with their clients as in the hypothetical scenario in which they do not trade with the clients (but nevertheless observe the clients' demand and thus make inferences about aggregate orderflow). Fixing a price $p_{1,i}$ and a quantity x_i, taking expectations of the utility defined in (2) as well as using the characterizations of the price p_2 and demand schedule τ in (3), (8) and (9), one gets

$$\mathbb{E}\left(U_i^d(x_i, p_{1,i}, p_2, \tau_i) \mid \text{clients demand } x_i, \text{ trade happens}\right)$$
$$= \mathbb{E}\left(-(\tau_i - x_i)^2 \tfrac{\gamma}{2} + x_i p_{1,i} - p_2 \tau_i\right)$$

for all $i \in \mathcal{D}$. If dealer i does not trade, but all other dealers do trade with their clients, dealer i expects the market price

$$p_2^{notrade}(x_1, ..., x_N) = \frac{-b}{aN}(\sum_{j=1, j\neq i}^{N} x_j) = \gamma(\bar{x} - \frac{x_i}{N}), \tag{12}$$

since dealer i optimally uses the demand schedule

$$\tau^{notrade}(p_2^{notrade}) := a p_2^{notrade} = \frac{N-2}{N-1}(\frac{x_i}{N} - \bar{x}),$$

while all other dealers use the demand schedule given by (3). Market clearing now implies (12). Analogously to the case in which the dealer trades with the clients, one gets

$$\mathbb{E}\left(U_i^d(x_i, p_{1,i}, p_2, \tau_i) \mid \text{clients demand } x_i, \text{ no trade happens}\right)$$
$$= \mathbb{E}\left(-(\tau_i^{notrade})^2 \tfrac{\gamma}{2} - p_2^{notrade} \tau_i^{notrade}\right)$$

Requiring that the utility that the dealer derives from trading and not trading, respectively, are the same and solving for $p_{1,i}$ gives

$$p_{1,i} = \mathbb{E}\left[\gamma \frac{\sum_{j=1}^{N} x_j (N-2) + x_i}{(N-1)N} \mid \text{clients demand } x_i\right]. \tag{13}$$

Using the distributional assumption on order flow (1) and the conditional expectation of a multivariate normal distribution gives

$$p_{1,i} = \gamma \frac{(N-2)\,\mathbf{1}'\Sigma_{i,i}^{-1}\Sigma_{*,i} + 1}{(N-1)N} x_i, \tag{14}$$

where $\mathbf{1}$ is an N-dimensional column vector and $\Sigma_{*,i}$ refers to the i-th column of the covariance matrix Σ. Using the expression $p_2 = -b/(aN)\sum_{j=1}^{N} x_j$, one has

$$\mathbb{E}\left[p_2 \mid \text{customer demands } x_i\right] = \frac{\gamma}{N}\mathbf{1}'\Sigma_{i,i}^{-1}\Sigma_{*,i}x_i. \tag{15}$$

Now, (14) and (15) imply

$$p_{1,i} = \frac{(N-2)}{(N-1)}\mathbb{E}\left[p_2 \mid \text{customer demands } x_i\right] + \gamma\frac{x_i}{(N-1)N}. \tag{16}$$

6.3 Client markups, dealer informedness, and price impact

One may define the markup as the difference between the quoted price (adjusted for the direction of the trade) and the quote of an uformed dealer of zero (since $p_2 = 0$ in expectation):

$$m_i := (\chi_{buy} - \chi_{sell})\left(\frac{(N-2)}{(N-1)}\mathbb{E}\left[p_2 \mid \text{net orderflow} = x_i\right] + \gamma\frac{x_i}{N(N-1)}\right), \tag{17}$$

where $\chi_{buy} = 1$ if a client buys, $\chi_{buy} = 0$ if a client sells and $\chi_{sell} = 1 - \chi_{buy}$. As the number of dealers gets large, we get

$$m_i \to (\chi_{buy} - \chi_{sell})\mathbb{E}\left[p_2 \mid \text{net orderflow} = x_i\right]. \tag{18}$$

A regression of dealer i's markups on the realized second-period prices (adjusted for the direction of the trade), i.e.

$$m_i = \beta_0 + \beta(\chi_{buy} - \chi_{sell})p_2 + \varepsilon, \tag{19}$$

is similar to the regression

$$\mathbb{E}\left[p_2 \mid \text{net orderflow} = x_i\right] = \beta_0' + \beta' p_2 + \varepsilon'. \tag{20}$$

We get from (19) to (20) using the limit in (18) and dropping the sign of the trade $\chi_{buy} - \chi_{sell}$ from both variables in the regression (19). In (20) we regress dealer i's conditional expectations of p_2 on the realized values of p_2 and we know that

$$\beta' \xrightarrow{p} R_i^2$$

as the sample becomes large, where R_i^2 is the r-squared statistic from the regression

$$p_2 = \beta_0 + \beta_1 x_i + \varepsilon,$$

i.e. a regression of the price p_2 on dealer i's order imbalance x_i.

If we run the regression described in (19), where m_i is replaced by the empirical estimates of the dealer's markups, and $(\chi_{buy} - \chi_{sell})p_2$ refers to the price impact measures, we expect that $\hat{\beta}$ is larger if R_i^2 is larger. We can obtain an estimate of R_i^2 from regressions of future price changes on the dealers' order imbalances.[10]

6.4 Endogenous dealer choice

We add two features to the model to capture the endogenous client-dealer choice.

First, we introduce an "arbitrageur" who is endowed with superior information about futures prices. The arbitrageur has access to the interdealer market and can request a quote from one dealer in the first trading round. In the first period, the arbitrageur and other customers of the contacted dealer have to pay the same price.

Second, dealers will have a different reservation utility, when giving quotes to their customers and may refuse to offer a quote if no quote can be found at which they break even.

All other assumptions on dealers and their clients will remain the same as before.

The arbitrageur can send a request for quote (RFQ) to one dealer in the first period, specifying some (exogenously given) quantity $\alpha > 0$ that the arbitrageur wants to trade. The RFQ does not specify whether the arbitrageur wants to buy or sell. Thus, the dealer responds with a bid-ask spread. After the arbitrageur has communicated whether she wants to buy or sell, the dealer charges the same price to the arbitrageur and to the uninformed customers. The uninformed customers are, as before, happy to trade at any price. Conditional on having traded with a dealer in period 1, the arbitrageur offloads the entire quantity α in the interdealer market in period 2.

In the interdealer market, the optimal demand schedule of a dealer who only traded with uninformed customers is still determined by (5). If, additionally, dealer i traded with the arbitrageur, then her inventory in period 2 is either $x_i - \alpha$ or $x_i + \alpha$, depending on whether the informed trader bought or sold. Thus, the first order condition for characterizing the maximum of (2) using the different initial inventory implies

[10] Since $\mathbb{E}(p_2) = 0$, we get from the law of iterated expectations that $\beta_0' \xrightarrow{p} 0$ in large samples. Using this limit and mutiplying both sides in (20) by $\chi_{buy} - \chi_{sell}$ gives $\beta_0 \xrightarrow{p} 0$ and $\beta \xrightarrow{p} \beta'$ in large samples, since $\chi_{buy} - \chi_{sell}$ is uncorrelated with the error in (20) (due to symmetry of the normal distribution, signed order flow predicts signed price changes. The sign has no additional predictive power). We thus get $\beta \xrightarrow{p} R_i^2$ in large samples.

$$\tau_i^b := \frac{\gamma(x_i + \alpha) - p_2}{\gamma + \lambda} \tag{21}$$

if the arbitrageur bought in period 1 and

$$\tau_i^s := \frac{\gamma(x_i - \alpha) - p_2}{\gamma + \lambda} \tag{22}$$

if the arbitrageur sold.

Since the demand of the arbitrageur is price insensitive, the price impact λ is still given by (4), where a, and b are defined in (8) and (9). If the arbitrageur decides to trade with dealer i, the market clearing condition in the interdealer market becomes

$$0 = \begin{cases} \sum_{j \neq i}^{N} \tau_j + \tau_i^b - \alpha & \text{if arbitrageur bought in period 1,} \\ \sum_{j \neq i}^{N} \tau_j + \tau_i^s + \alpha & \text{if arbitrageur sold in period 1.} \end{cases} \tag{23}$$

In equation (23), the quantity α enters with a negative sign if the arbitrageur bought in period 1, since, in that case, the arbitrageur will sell the same quantity in period 2. Solving (23) for p_2, using , (4), (5), (8), (9), (21) and (22) gives

$$p_2^b := \frac{\gamma \sum_{j=1}^{N} x_j}{N} - \frac{\gamma\alpha}{N(N-2)} \tag{24}$$

if the arbitrageur bought in period and

$$p_2^s := \frac{\gamma \sum_{j=1}^{N} x_j}{N} + \frac{\gamma\alpha}{N(N-2)} \tag{25}$$

if the arbitrageur sold in period 1.

The prices a dealer quotes when contacted by the arbitrageur are determined as follows.

- *ask price:* At the ask, the dealer is indifferent between selling α to the arbitrageur and trading with the uninformed customers and not trading in period 1, assuming that the arbitrageur will the buy α from another dealer, that the uninformed customers will also trade with another dealer and that the price in the interdealer market is *higher* than the ask price.

- *bid price:* At the bid, the dealer is indifferent between buying α from the arbitrageur and trading with the uninformed customers and not trading in period 1, assuming that the arbitrageur will sell α to another dealer, that the uninformed customers will also trade with another dealer, and that the price in the interdealer market is *lower* than the bid price.

Formally, this means

$$0 = \mathbb{E}\left[\left(-(\tau_i^b - x_i - \alpha)^2 \cdot \gamma/2 + (x_i + \alpha) \times ask - p_2^b \tau_i\right) - \left(-(\tau^{notrade})^2 \cdot \gamma/2 - p2^b \tau^{notrade}\right) \mid x_i, p_2 > ask\right]$$

Solving the last equation for ask gives

$$ask = \frac{N}{N+1}\mathbb{E}\left[p_2^b \mid p_2 > ask, x_i\right] + \frac{\gamma x_i}{(N-1)^2}. \tag{26}$$

Analogously, one gets

$$bid = \frac{N}{N+1}\mathbb{E}\left[p_2^s \mid p_2 < bid, x_i\right] + \frac{\gamma x_i}{(N-1)^2}. \tag{27}$$

For the sake of clarity, we now neglect any influence of the dealer's uninformed orderflow x_i on markups, i.e. we consider the limiting case for $\sigma_i \to 0$ and $\alpha \to 0$. Then the arbitrageur, who knows p_2 will buy from the dealer if and only if

$$p_2 > \frac{N}{N+1}\mathbb{E}\left[p_2 \mid x_i, p_2 > ask\right],$$

where p_2 is again determined as in Section 6.2.

Analogously, the arbitrageur will sell the asset to the dealer if and only if

$$p_2 < \frac{N}{N+1}\mathbb{E}\left[p_2^s \mid x_i, p_2 < bid\right].$$

The ask price converges in probability to $\frac{N}{N+1}p_2$ if $(corr(x_i, p_2))^2 \to 1$. Thus, the probability that the arbitrageur buys if $p_2 > 0$ goes to one if $(corr(x_i, p_2))^2 \to 1$.

Analogously, one can show that the probability that the arbitrageur will sell to the dealer goes to one if $p_2 < 0$ and $(corr(x_i, p_2))^2 \to 1$. Since $p_2 \neq 0$ almost surely, the arbitrageur will trade with a probability that is arbitrarily close to 1 if contacting a sufficiently informed dealer.

On the other hand, consider the case in which $corr(x_i, p_2) = 0$. Then, it can be shown that the dealer will charge a fixed positive bid-ask spread with midpoint zero and the arbitrageur will not trade in the (positive-probability) event that p_2 falls into that spread. Proposition B.1 formalizes and proves these statements.

6.5 Discussion

The model presented above describes a mechanism according to which dealers may include expectations about future prices because the dealers' reservation price, i.e. the price at which dealers are indifferent between trading and not trading, changes.

In actual markets, dealers may not necessarily quote their own reservation price, but may add a markup that depends on the clients bargaining power. The model does not analyse those additional markups. However, in empirical analyses it may be necessary to control for determinants of those markups in order to determine the proper effect of expected price changes on dealers' quotes.

7 Conclusion

Extending previous work on informed trading in FX markets, we document heterogeneity in the traders' informedness, as measured by their price impact. We also show that informedness is persistent, i.e. traders with a high price impact in one period are likely to have a high price impact in another period. Moreover, we present evidence suggesting that informed traders are more likely to trade with informed dealers. These findings are consistent with recent theory papers (Lee and Wang (2019)), Glode and Opp (2019)) that argue that non-anonymous OTC markets with a two-tiered market structure with "intermediation chains" may help alleviate asymmetric information frictions. Our findings are also consistent with Evans and Lyons (2002) and suggest that information frictions are also prevalent in a large market such as the Foreign-Exchange market, even though its underlying fundamentals are traditionally largely thought of as driven by systematic macro-economic risks. This has implications for the regulation of financial markets. If OTC markets with their two-tiered structure exist largely due to the rent-seeking behavior of a small set of large dealers who find ways to stymie competition, then regulators should mandate trading on exchanges, which would reduce excessive markups and, by lowering trading costs, lead to more efficient trade.[11] If, on the other hand, OTC markets help alleviate significant financial frictions, then it is less clear that all assets should be traded on anonymous centralized exchanges. The results in this paper offer some support in favor of this second view, and thus suggest that caution may be warranted when regulating financial market structure.

[11] Convincing evidence that channeling trades of smaller less sophisticated traders onto electronic trading platforms indeed leads to less price dispersion and more competitive prices is given in Hau et al. (2019). See also Duffie (2012).

Appendix

A Additional description of the data

Figure A.1: Histogram of maturity dates. This figure shows the frequency with which contracts of the maturities with values on the horizontal axis are traded. The sample period is May 2018 to April 2019. Maturity is expressed in days.

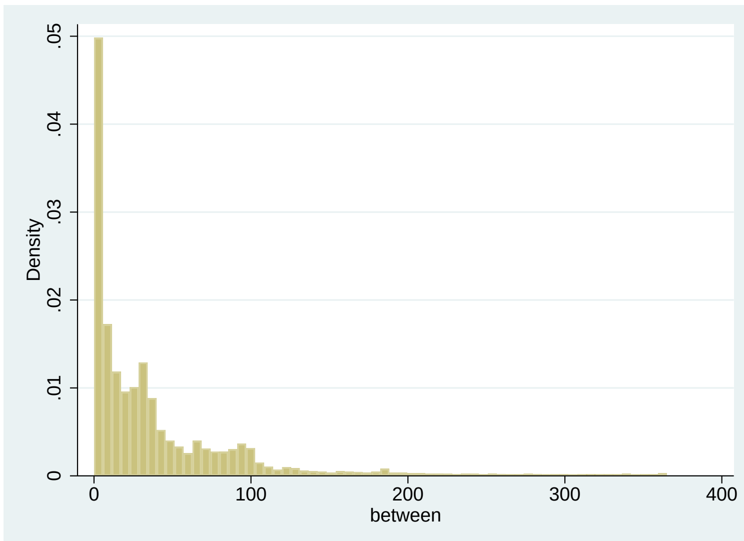

Table A.1: Trades in the C2C market. This table shows the number of trades that are executed between the dealers' customers directly in the sample period from May 2018 to April 2019. Stars signal that values have been omitted due to confidentiality concerns.

	Non-financial	Fund	Ins. & Pen.	Central Bank	Government	Empty
Non-financial	7,115					
Fund	247,269	165,206				
Insurance & Pension	312	991	8			
Central Bank	0	0	0	0		
Government	279	509	7	0	0	
Empty	14,351	125,155	891	98	92	*

Table A.2: Average notional in the C2C market. This table shows the average notional in EUR of trades that are executed between the dealers' customers directly in the sample period from May 2018 to April 2019.

	Non-financial	Fund	Ins. & Pen.	Central Bank	Government	Empty
Non-financial	892,935					
Fund	125,273	3,454,508				
Insurance & Pension	333,659	6,798,043	44,725,209			
Central Bank	0	0	0	0		
Government	11,874,660	4,830,630	10,092,969	0	0	
Empty	6,711,257	573,142	67,393,467	25,702,818	110,600,000	34,234

Table A.3: Average notional and number of trades in the D2D market. This table shows the number of trades and the average notional in EUR of trades that are executed between the dealers' customers directly in the sample period from May 2018 to April 2019. Stars signal that values have been omitted due to confidentiality concerns.

	number of trades		average notional volume	
	G16	Bank	G16	Bank
G16	995,490		98,028,476	
Bank	383,859	*	90,904,190	40,811,433

Figure A.2: Markups. This figure shows the average markups across all traders over the sample period.

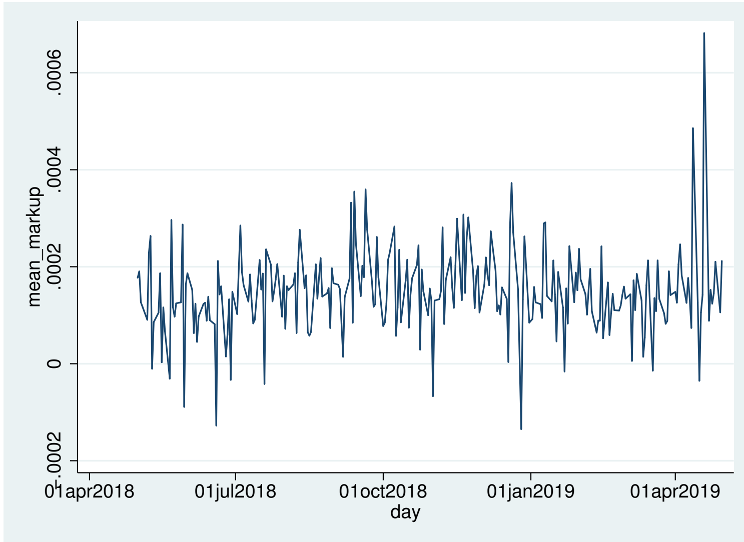

Figure A.3: Histogram of average trades per day. This figure shows the distribution of the trader's average number of trades per day conditional on trading. All traders that trade 12 time or more on average per day when trading have been allocated to the last group, i.e. to the 11-trades group.

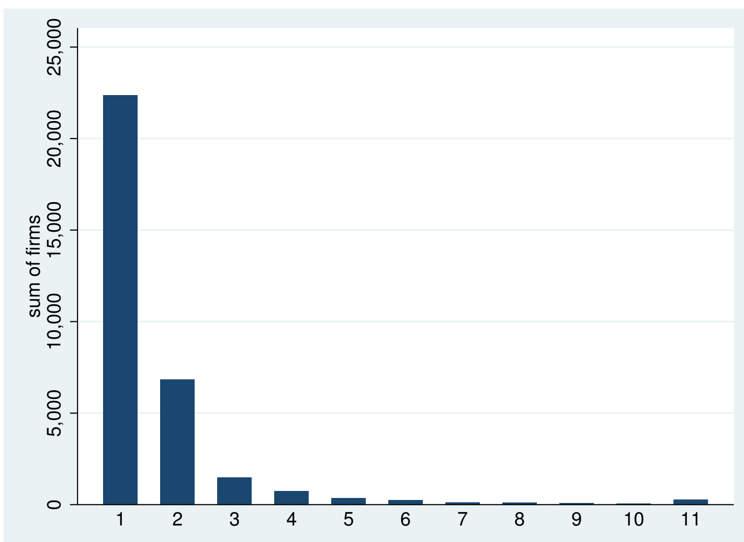

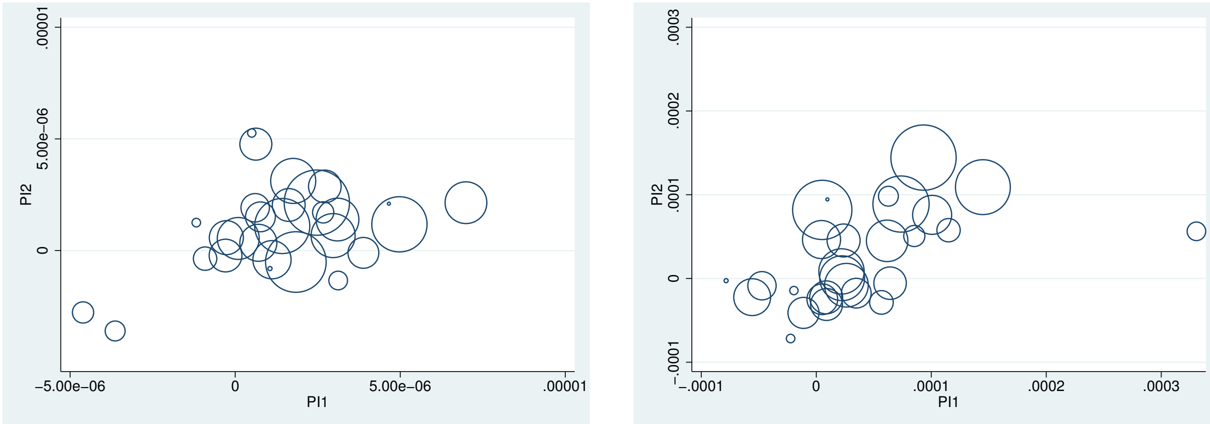

Panel A (1 minute).

Panel B (1 day).

Figure A.4: Persistence of price impact and the size of markups. In Panel A, a group's average 1-minute price impact in the second half of the sample is plotted against its 1-minute price impact in the first half of the sample. In Panel B, a group's average 1-day price impact in the second half of the sample is plotted against its 1-day price impact in the first half of the sample. Larger circles correspond to groups in which trades are associated with higher markups on average in the first half of the sample.

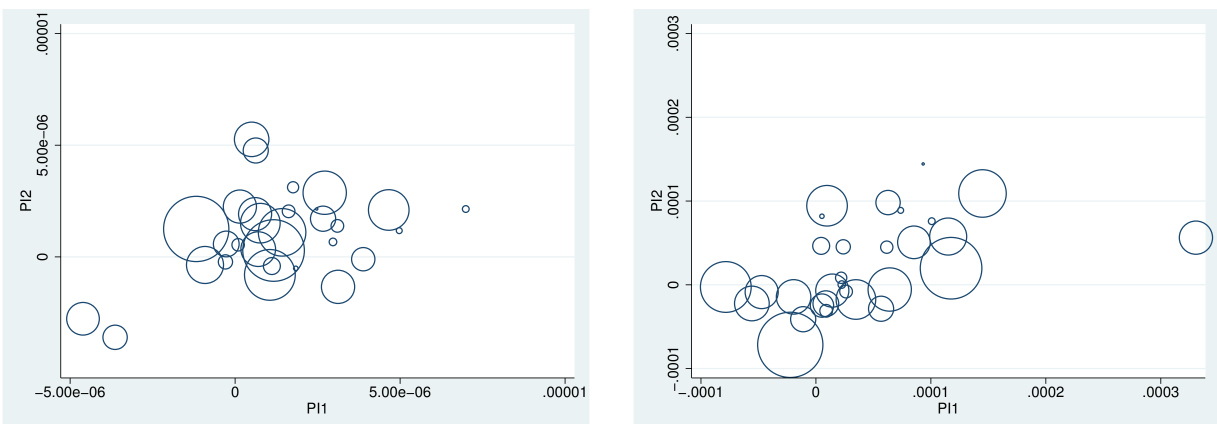

Panel A (1 minute).

Panel B (1 day).

Figure A.5: Persistence of price impact and the size of total notional traded. In both panels, a group's average 1-day price impact in the second half of the sample is plotted against its 1-day price impact in the first half of the sample. In Panel A, only HFT market participants are considered and in B, only non-HFT market participants are considered. Larger circles correspond to groups in which traders generate a higher total notional volume (EUR) in the first half of the sample.

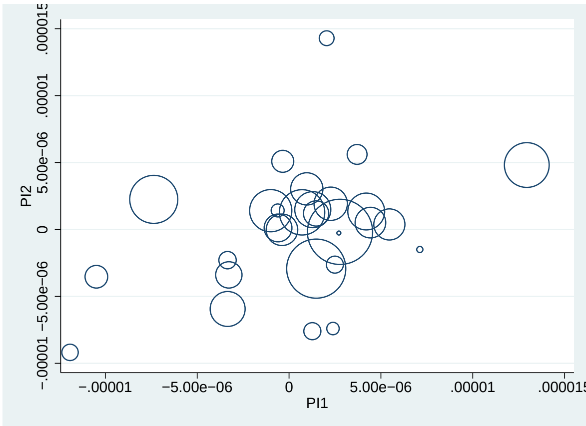

Panel A (HFT).

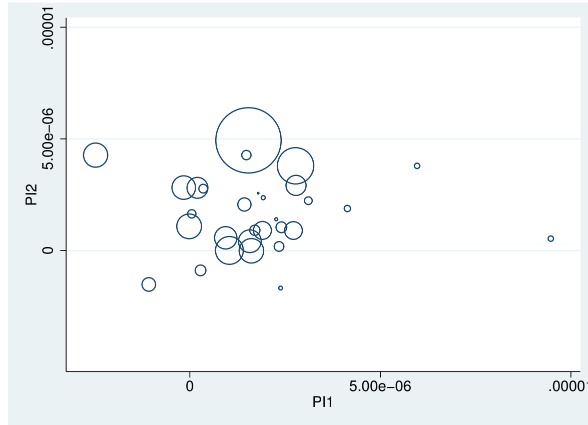

Panel B (non-HFT).

Figure A.6: Persistence of 1-minute price impact and the size of total notional traded for HFT and non-HFT market participants. In both panels, a group's average 1-minute price impact in the second half of the sample is plotted against its 1-minute price impact in the first half of the sample. In Panel A, only HFT market participants are considered and in B, only non-HFT market participants are considered. Larger circles correspond to groups in which traders generate a higher total notional volume (EUR) in the first half of the sample.

Figure A.8: Histogram of the dealers' informedness.

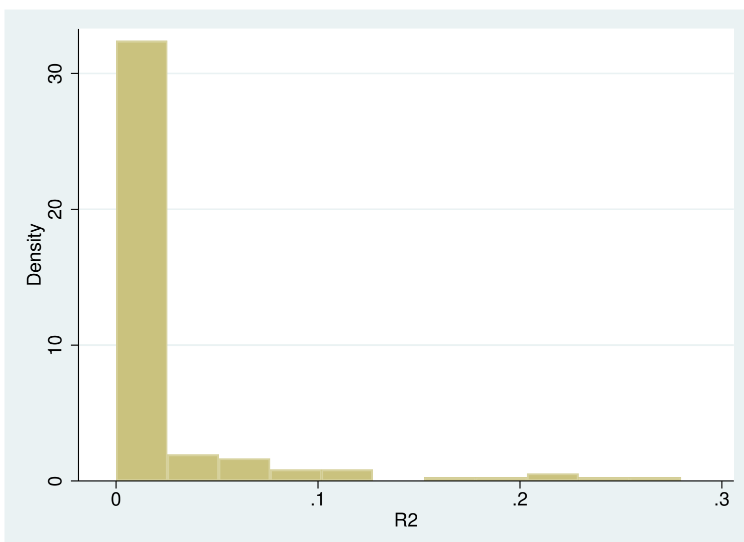

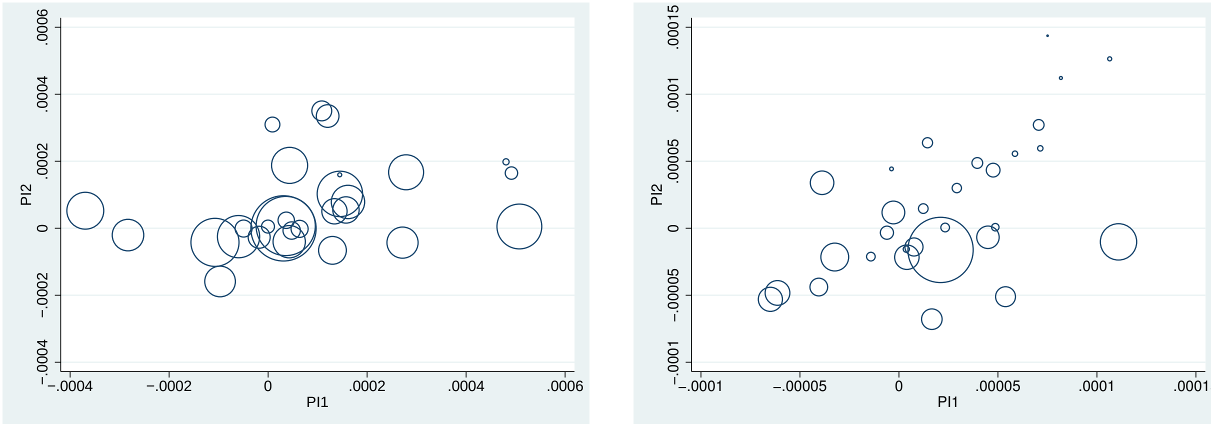

Panel A (HFT). Panel B (non-HFT).

Figure A.7: Persistence of 1-day price impact and the size of total notional traded for HFT and non-HFT market participants. This figure a group's average 1-minute price impact in the second half of the sample against its 1-minute price impact in the first half of the sample. Larger dots correspond to groups in which traders generate a higher notional volume (EUR) per trade in the first half of the sample.

Table A.4: Permanent price impact of traders. This table shows the average changes in the benchmark rates for the 1-minute and the 1-day horizon, the corresponding standard deviations as well as t-stats for each trader type. The standard deviations refer to single observations as opposed to the mean estimates.

	1-day impact			1-minute impact		
trader type	mean	std. dev.	t-stat	mean	std. dev.	t-stat
CENTRAL BANK	-0.0000133	0.0048531	-0.21	0.00000326	0.0002118	1.16
EMPTY	0.0000224	0.0046875	3.39	0.00000106	0.0002322	3.23
FUND	0.0000189	0.0046789	4.99	0.0000011	0.0002296	5.91
GOVERNMENT	-0.00000701	0.004862	-0.24	0.00000564	0.0002611	3.57
INSURANCE & PENSION	-0.0000763	0.0049035	-4.71	0.00000216	0.0002271	2.88
NON-FINANCIAL	0.0001174	0.0048899	19.03	0.00000133	0.0002408	4.38

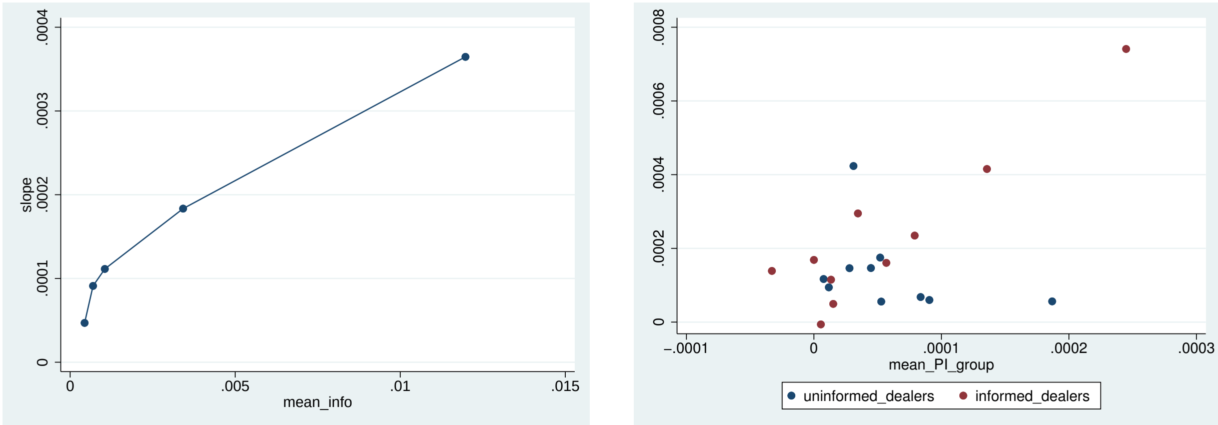

Panel A. Panel B.

Figure A.9: Informedness and sensitivity with respect to price impact. After having sorted trades first into quintiles based on dealer informedness and then having sorted trades within each quintile into two groups based on 1-day price impact, this figure shows the difference in the average markups of the group with a high 1-day price impact and the group with a low 1-day price impact for the 5 quintiles. The difference in markups is called slope. In Panel B, After having sorted trades first into terciles based on dealer informedness and then having sorted trades within each tercile into 10 groups based number of trades (holding the number roughly constant within each group), this figure shows average markups for each group an the average 1-day price impact of the trades of each group. Red dots (informed_dealers) refer to trades where the dealer informedness is in the highest tercile, while blue dots (uninformed_dealers) refer to trades where the dealer informedness is in the lowest tercile.

Table A.5: Price impact of dealers' clients. This table shows the average clients price impact for the 1-minute and the 1-day horizon as well as the corresponding standard deviations aggregated for each dealer type.

	1-day impact			1-minute impact		
trader type	mean	std. dev.	t-stat	mean	std. dev.	t-stat
BANK	0.0000706	0.004668	12.54	0.0000011	0.0002506	3.74
G16	0.0000278	0.004762	8.45	0.0000013	0.0002267	8.03

Table A.6: Markups for traders. This table shows the average markup for each trader type as well as the corresponding standard deviations. Stars indicate that values are not reported due to confidentiality concerns.

trader type	mean	standard deviation
CENTRAL BANK	*	0.0009116
EMPTY	0.0001678	0.0042364
FUND	0.0000754	0.0020724
GOVERNMENT	0.0001240	0.0014073
INSURANCE & PENSION	0.0000091	0.0015384
NON-FINANCIAL	0.0003284	0.0055631

Table A.7: Markups from dealers. This table shows the average markup for each dealer type as well as the corresponding standard deviations.

trader type	mean	standard deviation
BANK	0.000317	0.004075
G16	0.000092	0.003378

Table A.8: Informedness and price impact. This table shows the average price impact of clients whose dealers have informedness above the median compared to analogous statistics when dealers' informedness is below the median. Informedness is defined as in Section 4. Besides average 1-day and 1-minute price impact, the table shows the total number of trades executed by each group of dealers as well as the number of dealers that belongs to each group.

	1-day price impact	1-minute impact	trades	number of dealers
uninformed	0.000036600	0.000000538	1447095	34
informed	0.000041900	0.000001960	1359241	108

Table A.9: Connectedness in D2D market and price impact. This table shows the average price impact of clients whose dealers have more connections in the D2D market than the median compared to analogous statistics when dealers have fewer connections than the median. Besides average 1-day and 1-minute price impact, the table shows the total number of trades executed by each group of dealers as well as the number of dealers that belongs to each group.

	1-day price impact	1-minute impact	trades	number of dealers
unconnected	.000039	0.0000013200	1,508,747	125
connected	.000044	0.0000009440	1,237,597	7

Table A.10: Connectedness in D2C market and price impact. This table shows the average price impact of clients whose dealers have more connections in the D2C market than the median compared to analogous statistics when dealers have fewer connections than the median. Besides average 1-day and 1-minute price impact, the table shows the total number of trades executed by each group of dealers as well as the number of dealers that belongs to each group.

	1-day price impact	1-minute impact	trades	number of dealers
unconnected	0.000048	0.00000117	1,595,402	137
connected	.0000276	0.00000131	1,210,934	5

Table A.11: Probability of being an informed dealer. We run the regression

$$R^2 = BX + \varepsilon,$$

where R^2 is the informedness measure described in the text. This table shows the coefficient estimates B for various explanatory variables X as well as robust standard errors. In columns 1 to 3 we focus on dealers executing more than 0.5% of the notional volume of the entire D2C market. In columns 4-6 we focus on dealers executing more than 2.5% of the notional volume (EUR) of the entire D2C market.

	> 0.5% notional volume			> 2.5% notional volume		
	(1)	(2)	(3)	(4)	(5)	(6)
total notional traded notional	-0.96		0.04	-0.60		-2.90***
	(1.05)		(1.60)	(0.91)		(0.65)
# trades dealer		-0.10**	-0.01		-0.12**	-0.22***
		(0.05)	(0.07)		(0.04)	(0.05)
D2D counterparties			-0.00			
			(0.00)			
D2C counterparties			-0.00			
			(0.00)			
Constant	0.00**	0.00***	0.01**	0.00*	0.00***	0.01***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
N	20.00	20.00	20.00	11.00	11.00	11.00

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.12: Markups and price impact in different subsamples with fixed effects Using dealer-client fixed effects, we run the regression

$$markup_{it} = BX + \varepsilon,$$

where ε_{it} is an error term, X are explanatory variables specified in the table and $markup_{it}$ is the markup that trader i has to pay at time t. We report coefficients B and standard errors that are clustered at the dealer level. In column 1, we focus on HFT market participans, and in column 2 we focus on non-HFT market participants. In column 3 we exclude all traders labelled as neg-PI traders. In column 4 we focus only on those traders. In column 5 we focus on the trades for which volatility is below the median across all trades and in column 6 we focus on trades for which volatility is above the median across all trades. We excluded all trades with markups below -2% and above 3%. Order imbalance has been divided by 10^6.

	(1)	(2)	(3)	(4)	(5)	(6)
	HFT	non-HFT	no low-PI	low-PI	low vol	high vol
real. 1-day impact × info dummy	0.006	0.011**	0.008**	0.019***	0.008**	0.011*
	(0.004)	(0.005)	(0.004)	(0.007)	(0.003)	(0.006)
realized 1-day impact	0.009***	0.009***	0.009***	0.007***	0.008***	0.009***
	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)
realized 1-minute impact min	0.150***	0.061**	0.083***	0.148*	0.103**	0.085***
	(0.049)	(0.029)	(0.026)	(0.075)	(0.051)	(0.029)
market conditions:						
volatility	1.126	0.444*	0.870**	0.107	0.553	0.685
	(0.699)	(0.235)	(0.409)	(0.434)	(0.723)	(0.545)
smart average 1-day impact	0.004	-0.000	0.003	-0.005	0.007	-0.004
	(0.009)	(0.007)	(0.006)	(0.012)	(0.005)	(0.011)
varying trader characteristics:						
log(traders' monthly counterparties)	0.014	-0.027**	-0.005	-0.077***	0.002	-0.017
	(0.046)	(0.011)	(0.022)	(0.028)	(0.028)	(0.018)
log(trader's monthly trades)	0.032**	0.023**	0.032***	-0.008	0.011	0.040***
	(0.012)	(0.010)	(0.007)	(0.042)	(0.012)	(0.011)

(To be continued)

Table A.12-Continued.

	(1) HFT	(2) non-HFT	(3) no low-PI	(4) low-PI	(5) low vol	(6) high vol
varying dealer characteristics:						
dealer's signed OI	0.048	0.020	0.038	0.015	0.044	0.026
	(0.041)	(0.024)	(0.034)	(0.026)	(0.040)	(0.025)
Constant	-0.000	0.000***	0.000	0.000	0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
N	961931	1780807	2286548	456190	1384874	1357864

$p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.13: Probability of trading with connected dealers. We run the regression

$$connected = BX + \varepsilon,$$

where ε_{it} is an error term, X are explanatory variables specified in the table and *connected* is the the connectedness dummy discussed in Section 4. We report coefficients B and standard errors that are clustered at the trader level. The dataset includes D2C trades. We excluded all trades with markups below -2% and above 3%.

	excluding HFT			HFT only		
	(1)	(2)	(3)	(4)	(5)	(6)
avg. 1-day impact	-2.26	-2.11	-3.09	18.18	19.86	-11.90
	(3.57)	(3.56)	(3.41)	(73.43)	(72.51)	(70.97)
avg. 1-min impact		33.42	-29.67		320.21	-53.76
		(77.69)	(73.73)		(2227.70)	(2190.89)
realized 1-day impact		0.03	0.03		-0.09	-0.23*
		(0.05)	(0.05)		(0.13)	(0.12)
realized 1-min impact		-2.38**	-2.24**		-0.28	0.28
		(1.04)	(1.04)		(1.85)	(2.07)
log(counterparties)			68.15***			-69.73
			(9.05)			(50.16)
log(monthly trades)			-48.87***			115.78***
			(6.00)			(44.19)
volatility			-929.89***			65.94
			(120.59)			(597.69)
Constant	0.42***	0.42***	0.58***	0.51***	0.52***	-0.21
	(0.01)	(0.01)	(0.01)	(0.07)	(0.07)	(0.31)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B Technical details

This appendix contains technical results related to the model extension in section 6.4. First, it is shown that there are unique solutions to the fixed-point problems in (26) and (27). To this end, a result on truncated normal random variables is useful.

Definition 1. *A* truncated normal random variable *with parameters* μ *and* σ *and threshold* a *has the distribution of a normal random variable* Y *with mean* μ *and variance* σ^2 *conditional on* $Y > a$. *Thus,* X *has has the density*

$$\frac{\mathrm{d}\mathbb{P}(\{X < x\})}{\mathrm{d}x} = \begin{cases} \frac{\frac{1}{\sigma}\phi\left(\frac{x-\mu}{\sigma}\right)}{1-\Phi\left(\frac{x-\mu}{\sigma}\right)} & \textit{if } x \geq a, \\ 0 & \textit{if } x < a, \end{cases} \tag{B.1}$$

where ϕ *denotes the density of a standard normal random variable and* Φ *denotes the distribution function of a standard normal random variable.*

Lemma 1. *Let* X *be a truncated normal random variable with parameters* μ, σ *and threshold* a. *Let*

$$\kappa := \frac{a-\mu}{\sigma}. \tag{B.2}$$

Then,

$$\mathbb{E}(X) = \mu + \sigma\lambda(\kappa) \tag{B.3}$$

and

$$\mathbb{V}(X) = \sigma^2 \left[1 - \lambda(\kappa)(\lambda(\kappa) - \kappa)\right], \tag{B.4}$$

where λ *denotes the hazard rate function of a standard normal random variable, i.e.*

$$\lambda(x) := \frac{\phi(x)}{1-\Phi(x)}. \tag{B.5}$$

Proof. Using the density in (B.1), the moment generating function of X is given by

$$
\begin{aligned}
M(t) &:= \mathbb{E}\left(e^{tX}\right) \\
&= \frac{1}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)} \frac{1}{\sigma\sqrt{2\pi}} \int_a^\infty e^{ts} e^{-\frac{(s-\mu)^2}{2\sigma^2}} \, \mathrm{d}s \\
&= \frac{1}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)} \frac{1}{\sigma\sqrt{2\pi}} e^{\mu t+\sigma^2 t^2/2} \int_a^\infty e^{-\frac{(s-\mu-\sigma^2 t)^2}{2\sigma^2}} \, \mathrm{d}s \\
&= e^{\mu t+\sigma^2 t^2/2} \frac{1-\Phi\left(\frac{a-\mu}{\sigma}-\sigma t\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)}
\end{aligned}
$$

The first two derivatives are given by

$$
M'(t) = (\mu+\sigma^2 t) e^{\mu t+\sigma^2 t^2/2} \frac{1-\Phi\left(\frac{a-\mu}{\sigma}-\sigma t\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)} + e^{\mu t+\sigma^2 t^2/2} \frac{\sigma\phi\left(\frac{a-\mu}{\sigma}-\sigma t\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)}
$$

and (using $\phi'(x) = -x\phi(x)$)

$$
\begin{aligned}
M''(t) &= \left(\sigma^2 + (\mu+\sigma^2 t)^2\right) e^{\mu t+\sigma^2 t^2/2} \frac{1-\Phi\left(\frac{a-\mu}{\sigma}-\sigma t\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)} + 2(\mu+\sigma^2 t) e^{\mu t+\sigma^2 t^2/2} \frac{\sigma\phi\left(\frac{a-\mu}{\sigma}-\sigma t\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)} \\
&\quad + \sigma^2 \left(\frac{a-\mu}{\sigma}-\sigma t\right) e^{\mu t+\sigma^2 t^2/2} \frac{\phi\left(\frac{a-\mu}{\sigma}-\sigma t\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)}.
\end{aligned}
$$

Evaluating those derivatives at zero gives

$$
\mathbb{E}(X) = M'(0) = \mu + \sigma \frac{\phi\left(\frac{a-\mu}{\sigma}\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)}
$$

and

$$
\mathbb{E}(X^2) = M''(0) = \mu^2 + \sigma^2 + 2\mu\sigma \frac{\phi\left(\frac{a-\mu}{\sigma}\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)} + \sigma^2 \left(\frac{a-\mu}{\sigma}\right) \frac{\phi\left(\frac{a-\mu}{\sigma}\right)}{1-\Phi\left(\frac{a-\mu}{\sigma}\right)}.
$$

Using (B.2), (B.5) and $\mathbb{V}(X) = \mathbb{E}(X^2) - \mathbb{E}(X)^2$, one now gets (B.4).

□

The results in Lemma 1 can be used to show the existence and uniqueness of a bid ask spread that a dealer is willing to quote once contacted by an arbitrageur.

Lemma 2. *For any given customer demand x_i and a given covariance matrix Σ of the demanded quantities x_j, $j = 1, ..., N$, there are a unique ask $\in \mathbb{R}$ and a unique bid $\in \mathbb{R}$ such that (26) and (27) hold.*

Proof. This proof focuses on a solution to (26), since the corresponding statement for (27) is shown analogously. For a fixed demand of uninformed customers x_i, consider the function $f : \mathbb{R} \to \mathbb{R}$ defined by

$$f(t) := \frac{N}{N+1}\mathbb{E}\left[p_2|x_i, p_2 > t\right] + \frac{\gamma x_i}{(N-1)^2}$$

such that (26) can be rewritten as $f(ask) = ask$. We also have

$$\mathbb{E}\left[p_2|x_i, p_2 > t\right] = \frac{\int_t^\infty s\rho(s)\,\mathrm{d}s}{\int_t^\infty \rho(s)\,\mathrm{d}s},$$

where ρ is the density of p_2 with respect to the Lebesgue measure of p_2 conditional on the realization of x_2. Due to (24), the joint normality and analogously to the reasoning in the text before (15), this means (by the normal projection theorem) that μ is the density of a normal random variable with mean

$$\overline{\mu} := \frac{\gamma}{N}\mathbf{1}'\Sigma_{i,i}^{-1}\Sigma_{*,i}x_i - \frac{\gamma\alpha}{N(N-2)} \tag{B.6}$$

and variance

$$\overline{\sigma}^2 := \frac{\gamma^2}{N^2}\left[\mathbf{1}'\Sigma\mathbf{1} - \frac{\left(\mathbf{1}'\Sigma_{*,i}\right)^2}{\Sigma_{i,i}}\right]. \tag{B.7}$$

By (B.3) in Lemma 1, we have

$$\mathbb{E}\left(p_2|x_i, p_2 > K\right) = \overline{\mu} + \overline{\sigma}\lambda\left(\frac{K-\overline{\mu}}{\overline{\sigma}}\right). \tag{B.8}$$

By the Banach Fixed Point Theorem, there is a unique $ask \in \mathbb{R}$ with $ask = f(ask)$ if $|f'(x)| \leq M$ with $0 < M < 1$ for all $x \in \mathbb{R}$. The latter is shown in two steps:

1. The hazard rate function λ ist strictly increasing, i.e. $\lambda' > 0$.

2. The hazard rate function λ satisfies $\lambda' < 1$

Then, $\lambda'(x) \in (0,1)$ implies $f'(x) \in (0, \frac{N}{N+1})$. Thus, the claim follows by the Banach Fixed-Point Theorem.

To prove the first step, differentiate the hazard rate function. Using the facts that $\phi'(x) = -x\phi(x)$ (which can be verified by simple computation) and $\Phi'(x) = \phi(x)$ (which holds by definition), one gets

$$\begin{aligned}\lambda'(x) &= \frac{\phi'(x)(1-\Phi(x)) + \Phi'(x)\phi(x)}{(1-\Phi'(x))^2} \\ &= -x\frac{\phi(x)}{1-\Phi(x)} + \left(\frac{\phi(x)}{1-\Phi(x)}\right)^2 \\ &= \lambda(x)(\lambda(x) - x). \end{aligned} \tag{B.9}$$

Applying the formula (B.3) for the expectation of truncated normal random variables to a standard normal random variable X, one gets

$$\lambda(x) = \mathbb{E}(X|X > x) > x.$$

Since by (B.5) one has $\lambda > 0$, it now follows from (B.9) that $\lambda'(x) > 0$ for all $x \in \mathbb{R}$.

To prove the second step, note that the expression on the left-hand side of (B.4) must be positive for any $\kappa \in \mathbb{R}$. This implies that

$$\lambda(x)(\lambda(x) - x) < 1$$

for all $x \in \mathbb{R}$. Now it follows with (B.9) that $\lambda'(x) < 1$ for $x \in \mathbb{R}$.

□

In order to characterize the bid ask spread for different levels of dealer informedness, the following two auxiliary results are useful.

Lemma 3. *The expression $\sigma\lambda\left(\frac{x}{\sigma}\right)$, where λ is given by (B.5), is strictly monotone increasing in σ for all $x \in \mathbb{R}$ and all $\sigma > 0$.*

Proof. Let $x \in \mathbb{R}$ and $0 < \sigma_1 < \sigma_2$. As shown in the proof of Lemma 2, one has $0 < \lambda(y) > y$ and $\lambda'(y) < 1$ for all $y \in \mathbb{R}$. From these two fact it follows that the line $l : \left[\, 0, \frac{x}{\sigma_1}\, \right) \to \mathbb{R}$ with

$$l(y) := y\frac{\lambda(x/\sigma_1)}{x/\sigma_1}$$

lies strictly below λ on its domain. Using $y = x/\sigma_2$, this implies

$$\frac{\sigma_1}{\sigma_2}\lambda(x/\sigma_1) < \lambda(x/\sigma_2).$$

Multiplying both sides by σ_2 gives

$$\sigma_1\lambda\left(\frac{x}{\sigma_1}\right) < \sigma_2\lambda\left(\frac{x}{\sigma_2}\right).$$

□

We can now prove the claim that the arbitrageur is almost always able to trade with the dealer if the dealer is sufficiently informed. Moreover, the following results states the resulting bid-ask spreads for both informed and uninformed dealers and provides an expression for the probability with which an arbitrageur is able to trade with a completely uninformed dealer given that the dealer's order flow from uninformed clients is negligible.

Proposition B.1. *Let*

$$R^2 := \frac{\left(\mathbf{1}'\Sigma_{*,i}\right)^2}{\Sigma_{i,i}\mathbf{1}'\Sigma\mathbf{1}}.$$

be the squared correlation between dealer i's order flow from uninformed customers and the future price p_2*, let* $\Sigma_{i,i} \to 0$ *and let* $\alpha \to 0$*. Then,*

- *As* $R^2 \to 1$*, both the ask (bid) of dealer i converges in probability to* $\frac{N}{N+1}p_2$ *if* $p_2 > 0$ *(*$p_2 < 0$*) and to zero otherwise and the probability that the arbitrageur trades with the dealer goes to 1.*
- *If* $R^2 = 0$*, one has*

$$\frac{ask}{\gamma\sqrt{\mathbf{1}'\Sigma\mathbf{1}}/N} = S$$

and

$$\frac{bid}{\gamma\sqrt{\mathbf{1}'\Sigma\mathbf{1}}/N} = -S,$$

where S satisfies the fixed-point equation

$$S := \frac{N}{N+1}\lambda(S),$$

where λ *is given by* (B.5)*. In this case, the probability that the arbitrageur trades with the dealer is strictly below 1.*

In the following, the first and second bullet point are proved separately. Due to symmetry, only the expressions for the limit of the ask price will be derived explicitly.

Step1: Proof of the first bullet point.

With the ask price given by (26) and using (B.6), (B.7) and (B.8) from the proof of Lemma 2, one has

$$\begin{aligned} ask &\overset{p}{\to} \frac{N}{N+1}\mathbb{E}\left[p_2 \mid p_2 > ask, x_i\right] \\ &= \frac{N}{N+1}\left(\overline{\mu} + \overline{\sigma}\lambda\left(\frac{ask - \overline{\mu}}{\overline{\sigma}}\right)\right) \end{aligned} \tag{B.10}$$

as $\Sigma_{i,i} \to 0$. This convergence can be shown using Chebyshev's inequality. In the following, the limit in (B.10) will be examined for $R^2 \to 1$. First, $R^2 \to 1$ and (B.7) imply

$$\overline{\sigma} \to 0. \tag{B.11}$$

Moreover, for $\overline{\sigma} > 0$, it must always be the case that $ask > \frac{N}{N+1}\overline{\mu}$. Suppose that is not the case. Then one would get the contradiction

$$\frac{N}{N+1}\overline{\mu} \geq ask = \frac{N}{N+1}\left(\overline{\mu} + \overline{\sigma}\lambda\left(\frac{ask - \overline{\mu}}{\overline{\sigma}}\right)\right) > \frac{N}{N+1}\overline{\mu}.$$

Since, for any fixed $\overline{\mu} > 0$, the ask prices are bounded from below and are, by Lemma 3, strictly increasing in $\overline{\sigma}$, the prices must converge as $\overline{\sigma} \to 0$, holding $\overline{\mu}$ fixed.

Note that $0 < \lambda'(x) < 1$ implies that

$$\max\{\lambda(x) - x | x \geq 0\} = \lambda(0) = \sqrt{\frac{2}{\pi}}. \tag{B.12}$$

Suppose now, that $\lim_{\overline{\sigma}\to 0} ask \geq \overline{\mu}$. Then, using (B.12), one would obtain the contradiction

$$\lim_{\overline{\sigma}\to 0} ask = \lim_{\overline{\sigma}\to 0} \frac{N}{N+1}\left(\overline{\mu} + \overline{\sigma}\lambda\left(\frac{ask - \overline{\mu}}{\overline{\sigma}}\right)\right) \leq \lim_{\overline{\sigma}\to 0} \frac{N}{N+1}\left(\overline{\mu} + \overline{\sigma}\sqrt{\frac{2}{\pi}} + (ask - \overline{\mu})\right) = \frac{N}{N+1}\lim_{\overline{\sigma}\to 0} ask.$$

Thus,

$$\lim_{\overline{\sigma}\to 0} ask < \mu. \tag{B.13}$$

Since for any $x < \overline{\mu}$ and a standard normal random variable X, one has

$$\lim_{\overline{\sigma}\to 0} \lambda\left(\frac{x-\overline{\mu}}{\overline{\sigma}}\right) = \lim_{b\to -\infty} \mathbb{E}(X|X > b) = 0,$$

it must be the case that

$$\lim_{\overline{\sigma}\to 0} ask = \lim_{\overline{\sigma}\to 0} \frac{N}{N+1}\left(\overline{\mu} + \overline{\sigma}\lambda\left(\frac{ask-\overline{\mu}}{\overline{\sigma}}\right)\right) = \frac{N}{N+1}\overline{\mu}. \tag{B.14}$$

Suppose now that $\overline{\mu} < 0$. Because $ask - \overline{\mu} > -\frac{1}{N+1}\mu$ for any $\overline{\sigma} > 0$ due to the above reasoning, one has

$$\lim_{\overline{\sigma}\to 0} \overline{\sigma}\lambda\left(\frac{ask-\overline{\mu}}{\overline{\sigma}}\right) \to ask - \mu$$

because of (B.12). If now follows that the ask price has to satisfy

$$\lim_{\overline{\sigma}\to 0} ask = \frac{N}{N+1}(\overline{\mu} + \lim_{\overline{\sigma}\to 0} ask - \overline{\mu}),$$

i.e. $\lim_{\overline{\sigma}\to 0} ask = 0$.

Since $\overline{\mu} \xrightarrow{p} p_2$ as $\overline{\sigma} \to 0$ and $\alpha \to 0$, where p_2 is determined as in Section 6.2, the probability that $ask < p_2$ if $p_2 > 0$ goes to 1 as $\overline{\sigma} \to 1$. Analogously, the probability that $bid > p_2$ if $p_2 < 0$ goes to 1 as $\overline{\sigma} \to 1$. To conclude, the probability that the arbitrageur trades with the dealer goes to 1 as $\overline{\sigma} \to 1$, since $p_2 > 0$ or $p_2 < 0$ with probability 1.

Step2: Proof of the second bullet point.

As above, the limit in (B.10) is examined for $R^2 = 0$. The last two conditions and (B.7) imply

$$\overline{\sigma} \to \infty. \tag{B.15}$$

Moreover, since $R^2 = 0$, one always has $\overline{\mu} = 0$ as can be seen from (B.6), since $\mathbf{1}'\Sigma_{i,i}^{-1}\Sigma_{*,i} = 0$ in this case. Thus, (B.10) becomes

$$ask \xrightarrow{p} \frac{N}{N+1}\left(\overline{\sigma}\lambda\left(\frac{ask-\overline{\mu}}{\overline{\sigma}}\right)\right) \tag{B.16}$$

Instead of looking at the behavior of the ask prices directly, we consider the expression

$$K := \frac{ask}{\overline{\sigma}}.$$

If the ask price satisfies (B.16), then K satisfies

$$K = \frac{N}{N+1}\lambda(K). \tag{B.17}$$

Since $0 < \lambda'(x) < 1$ for all $x \in \mathbb{R}$, a solution K to (B.17) exists and is unique, which can be shown using the Banach Fixed-Point Theorem. Moreover, $K > 0$, since $\lambda(x) > 0$ for all $x \in \mathbb{R}$.

The probability that $p_2 > ask$ is given by $1 - \Phi(\frac{K}{\overline{\sigma}})$, where K satisfies (B.17) and Φ denotes the distribution function of a standard normal random variable. Since $K > 0$, this probability is less than $\frac{1}{2}$. Analogously, the probability that $p_2 < bid$ is less than $\frac{1}{2}$. Thus, the arbitrageur will trade with the uninformed dealer with a probability of less than 1. One obtains the expressions for the bid and ask prices in the statement of the proposition by using (B.7) to express $\overline{\sigma}$ in terms of the primitive parameters.

C The effect of measurement errors in the FX benchmark price on regression coefficient estimates

While most data are generally affected by measurement errors, a potential measurement error in the benchmark price (e.g. due to noisy quotes from dealers in the TRTH database or imprecise timestamps in the EMIR database) may pose a special problem since the benchmark price is used for both calculating price impact and markups. This appendix has three goals:

1. It is explored how errors in the benchmark price of the transactions affect the coefficient estimates in a regression of markups on price impact.

2. An argument is presented that, given the empirical results in the main text, the estimate of the impact of 1-day price impact on markups has, under plausible conditions, an upward bias of no more than 1 % of the coefficient estimate.

3. It is shown that even if errors in the benchmark price are large, this effect does, under plausible conditions, not affect estimates on how differently informed dealers respond to markups.

A potential bias in coefficient estimates: To address the first of the above points, let t be the reported time of a transaction that is observed in the dataset and let T be the point in time a fixed horizon after the reported transaction (e.g. $T - t$ is equal to one minute, one day, etc). Consider now the following five random variables:

$$m_t^o := \text{ observed FX benchmark price at time } t,$$

$$m_t^* := \text{ actual FX benchmark price at time } t$$

$$m_T^o := \text{ observed FX benchmark price at time } T,$$

$$m_T^* := \text{ actual FX benchmark price at time } T,$$

$$P := \text{price paid for the contract},$$

where the price of the contract is not affected by measurement errors. Moreover, let $sign(trade) := 1$ if the trader buys and $sign(trade) := -1$ if the trader sells. With the definitions given above, we can define the observed price impact, PI^o, the true price impact, PI^*, the observed markup, MU^o, and the true markup, MU^* as follows.

$$PI^o := sign(trade)(m_T^o - m_t^o),$$

$$PI^* := sign(trade)(m_T^* - m_t^*),$$

$$MU^o := sign(trade)(P - m_t^o),$$

$$MU^* := sign(trade)(P - m_t^*).$$

Using these definitions, we can express the observed variables in terms of the corresponding actual variables:

$$PI^o = PI^* + \underbrace{sign(trade)(m_t^* - m_t^o)}_{=:\varepsilon_1} + \underbrace{sign(trade)(m_T^o - m_T^*)}_{=:\varepsilon_2}, \tag{C.1}$$

$$MU^o = MU^* + \underbrace{sign(trade)(m_t^* - m_t^o)}_{=\varepsilon_1}, \tag{C.2}$$

where we also defined the error terms ε_1 and ε_2 that affect the observations of price impact and markup.

Consider now the univariate regression model

$$MU^* = \beta_0 + \beta_1 PI^* + \varepsilon, \tag{C.3}$$

where ε is an error term and $\beta_1 = \frac{Cov(MU^*, PI^*)}{Var(PI^*)}$. We can estimate β_1 by replacing actual values by observed values, but then, by (C.1) and (C.2), our estimate $\hat{\beta}_1$ becomes (in a large sample)

$$\hat{\beta}_1 = \frac{Cov(MU^o, PI^o)}{Var(PI^o)} = \frac{Cov(MU^* + \varepsilon_1, PI^* + \varepsilon_1 + \varepsilon_2)}{Var(PI^* + \varepsilon_1 + \varepsilon_2)}, \tag{C.4}$$

which is not necessarily equal to the true β_1.

Quantifying the bias: In the following, an upper bound for the bias derived above is stated for the case in which errors are uncorrelated with the true markup. It is also assumed that errors are uncorrelated with each other as well as with the true price impact. Let $\hat{\beta}_1^{minute}$ and $\hat{\beta}_1^{day}$ denote the estimates of the regression coefficient β_1 from (C.3), where PI^o stands for the observed 1-minute price impact or for the 1-day price impact, respectively. Note that the error ε_1 that affects those estimates is the same for the case with the 1-day price impact as for the case with the 1-minute price impact. The coefficient estimates $\hat{\beta}_1$ for various horizons are shown in Table C.1. As for the regressions in Section 4, trades by CCPs and central banks as well as trades with extreme markups have been excluded from the regressions shown in Table C.1.

Table C.1: Markups and price impact. This table shows coefficient estimates for an OLS regression of markups on realized values of the price impact for various horizons. Standard errors are clustered on the dealer level and shown in parentheses. Trades by central banks, CCPs as well as trades with markups smaller than 2% or greater than 3% have been excluded.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	all	non-HFT	HFT	all	all	all	all	all
1-min impact	0.081**	0.043	0.141***				0.085***	0.086**
	(0.032)	(0.029)	(0.049)				(0.032)	(0.034)
30-min impact				0.035***				0.026***
				(0.008)				(0.008)
1-day impact					0.016***		0.016***	0.006***
					(0.002)		(0.002)	(0.002)
5-day impact						0.013***		0.011***
						(0.002)		(0.001)
Constant	0.000***	0.000***	0.000**	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The estimates

If $Cov(MU^*, PI^*) \geq 0$, we get from (C.4) and and the correlations assumed above that

$$Var(\varepsilon_1) \leq \hat{\beta}_1^{minute} Var(PI^o). \tag{C.5}$$

Looking at Table C.1, we have $\hat{\beta}_1^{minute} \approx 0.1$. and $\hat{\beta}_1^{day} \approx 0.01$. Using standard deviations of the observed 1-minute price impact, $PI^{o,minute}$, and 1-day priceimpact, $PI^{o,day}$ from Tables A.4 and A.5, we get $Var(PI^{o,minute}) \approx (0.0002)^2$ and $Var(PI^{o,day}) \approx (0.005)^2$. Using these approximations, we get from (C.5) that

$$Var(\varepsilon_1) \lessapprox 4 \cdot 10^{-10}.$$

We also get

$$Cov(PI^{o,day}, MU^o) = \hat{\beta}_1^{day} \cdot Var(PI^{o,day}) \approx 3 \cdot 10^{-7}.$$

and therefore, considering (C.4), that

$$\begin{aligned}
\frac{\hat{\beta}_1^{day}}{\beta_1^{day}} &= \frac{Cov(M^o, PI^{o,day})}{Cov(M^o, PI^{o,day}) - Var(\varepsilon_1)} \frac{Var(PI^{o,day}) - Var(\varepsilon_1) - Var(\varepsilon_2)}{Var(PI^{o,day})} \\
&\leq \frac{Cov(M^o, PI^{o,day})}{Cov(M^o, PI^{o,day}) - Var(\varepsilon_1)} \\
&\approx \frac{3 \cdot 10^{-7}}{3 \cdot 10^{-7} - 4 \cdot 10^{-10}}.
\end{aligned}$$

The potential upward bias can therefore only be in the order of magnitude of 0.1% of the original coefficient estimate.

One can see in Table C.1 that the relationship between price impact and markups is very robust with respect to the horizon. Moreover, comparing model 1, model 3 and model 5, the coefficient estimates barely change, when including more variables. If the observations would be heavily affected by a common measurement error, we would expect that estimates would change more, since the independent variables would be heavily correlated. Only after including the 5-day price impact (model 6), the coefficient estimate for the 1-day price impact changes, which is plausibly explained by the correlation between 1-day and 5-day price changes.

Moreover, while 1-minute price impact is significantly related to the markups for non-HFT client, the same is not true for HFT clients. Unslesss the trades of HFT clients have on average a different error, this suggests that the strong result for non-HFT clients is not driven by the error.

No upward bias for the estimate for the interaction term: Consider the linear regression model

$$MU^* = \beta_0 + \beta_1 PI^* + \varepsilon + \beta_2 PI^* \mathbf{1}_{informed} + \varepsilon, \tag{C.6}$$

whith the same interpretion as (C.3) and $\mathbf{1}_{informed} = 1$ if the trade happens with a dealer that is classified as informed and $\mathbf{1}_{informed} = 0$ otherwise. As the sample becomes large, one has

$$\hat{\beta}_2 = \hat{\beta}_1^{informed} - \hat{\beta}_1^{uninformed},$$

where $\hat{\beta}_1^{informed}$ and $\hat{\beta}_1^{uninformed}$ are the estimates of β_1 from running the regression in (C.3) for informed dealers only and uninformed dealers only, respectively. Using the results from above and assuming that $Var(\varepsilon_1)$ as well as the variance of the price impact do not change across informed and uninformed dealers, we have

$$\hat{\beta}_2 = \frac{Cov(MI^*, PI^*|\text{ informed}) - Cov(MI^*, PI^*|\text{ uninformed})}{Var(PI^o)}$$

as the sample becomes large, whereas the true coefficient satisfies

$$\beta_2 = \frac{Cov(MI^*, PI^*|\text{ informed}) - Cov(MI^*, PI^*|\text{ uninformed})}{Var(PI^*)}.$$

Under the assumptions stated above, $Var(PI^o) \geq Var(PI^*)$. Thus, the positive estimate of β_2 may only be biased towards zero but not upwards if the estimate is positive.

References

Abad, J., I. Aldasoro, C. Aymanns, M. D'Errico, L. Fache, Rousova, P. Hoffmann, S. Langfield, M. Neychev, and T. Roukny (2016). Shedding light on dark markets: First insights from the new eu-wide otc derivatives dataset. *ECB working paper.*

Babus, A. and P. Kondor (2018). Trading and information diffusion in over-the-counter markets,. *Econometrica 86*, 1727–1769.

Bjønnes, G., N. Kathitziotis, and C. Osler (2017). Bid-ask spreads in otc markets. *working paper.*

Cochrane, J. H. (1999). Portfolio advice in a multifactor world. *Feceral Reserve Bank of Chicago - Economic Perspectives 23.*

Di Maggio, M., A. Kermani, and Z. Song (2017). The value of trading relations in turbulent times. *Journal of Financial Economics 124*, 266–284.

Duffie, D. (2012). *Dark Markets: Asset Pricing and Information Transmission in Over-the-Counter Markets.* Princeton University Press.

Evans, M. and R. Lyons (2002). Order flow and exchange rate dynamics. *Journal of Political Economy 110*, 170–180.

Evans, M. D. D. and R. K. Lyons (2005). Exchange rate fundamentals and order flow. *NBER working paper.*

Glode, V. and C. Opp (2016). Asymmetric information and intermediation chains. *American Economic Review 106*, 2699–2721.

Glode, V. and C. Opp (2019). On the Efficiency of Long Intermediation Chains. *Journal of Financial Intermediation 38*, 11–18.

Glosten, L. R. and P. R. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics 14*(1), 71–100.

Hagströmer, B. and A. J. Menkveld (2019). Information revelation in decentralized markets. *Journal of Finance*, forthcoming.

Hau, H., P. Hoffmann, S. Langfield, and Y. Timmer (2019). Discriminatory Pricing of Over-The-Counter Derivatives. *working paper.*

Hendershott, T. and A. Madhavan (2015). Click or Call? Auction versus Search in the Over-the-Counter Market. *The Journal of Finance 70*(1), 419–447.

Kim, O. and R. E. Verrecchia (1991). Market reaction to anticipated announcements. *Econometrica 30*(2), 273–309.

Kondor, P. and G. Pintér (2019). Clients' connections: Measuring the role of private information in decentralised markets. *working paper.*

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica 53*(6), 1315–1336.

Lee, T. and C. Wang (2019). Why trade over-the-counter? when investors want price discrimination. *working paper.*

Li, D. and N. Schürhoff (2019). Dealer Networks. *Journal of Finance 74*(1), 91–144.

Liu, Y., S. Vogel, and Y. Zhang (2018). Electronic trading in otc markets vs. centralized exchange. *working paper.*

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2017). Information Flows in Foreign Exchange Markets: Dissecting Customer Currency Trades. *The Journal of Finance 71*(2), 601–634.

Neklyudov, A., B. Hollifield, and C. Spatt (2017). Bid-Ask Spreads, Trading Networks, and the Pricing of Securitizations. *Review of Financial Studies 30*, 3048–3085.

Ranaldo, A. and F. Somogyi (2018). Heterogeneous information content of global fx trading. *working paper.*

Sambalaibat, B. (2018). Endogenous specialization and dealer networks. *working paper.*

Seppi, D. J. (1990). Equilibrium Block Trading and Asymmetric Information. *The Journal of Finance 45*(1), 73–94.

Wang, C. (2017). Core-Periphery Trading Networks. *working paper.*