BANK FOR INTERNATIONAL SETTLEMENTS

BIS Working Papers

No 563

Who supplies liquidity, how and when?

by Bruno Biais, Fany Declerck and Sophie Moinas

Monetary and Economic Department

May 2016

JEL classification: D82, D53, G01

Keywords: Liquidity, high-frequency trading, proprietary trading, adverse selection, electronic limit order book, short-term momentum, contrarian

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Foreword

The 14th BIS Annual Conference took place in Lucerne, Switzerland, on 26 June 2015. The event brought together a distinguished group of central bank Governors, leading academics and former public officials to exchange views on the topic "Towards 'a new normal' in financial markets?". The papers presented at the conference and the discussants' comments are released as *BIS Working Papers* nos 561 to 564.

BIS Papers no 84 contains the opening address by Jaime Caruana (General Manager, BIS), the keynote address by John Kay (London School of Economics) and remarks by Paul Tucker (Harvard Kennedy School).

Who supplies liquidity, how and when?

Bruno Biais, Fany Declerck and Sophie Moinas[1]

Abstract

Who provides liquidity in modern, electronic limit order book, markets? While agency trading can be constrained by conflicts of interest and information asymmetry between customers and traders, prop traders are likely to be less constrained and thus better positioned to carry inventory risk. Moreover, while slow traders' limit orders may be exposed to severe adverse selection, fast trading technology can improve traders' ability to monitor the market and avoid being picked off. To shed light on these points, we rely on unique data from Euronext and the AMF, the French financial markets regulator, enabling us to observe the connectivity of traders to the market, and whether they are proprietary traders. We find that proprietary traders, be they fast or slow, provide liquidity with contrarian marketable orders, thus helping the market absorb shocks, even during a crisis, and they earn profits while doing so. Moreover, fast traders provide liquidity by leaving limit orders in the book. Yet, only prop traders can do so without making losses. This suggests that technology is not enough to overcome adverse selection; monitoring incentives are also needed.

Keywords: Liquidity, high-frequency trading, proprietary trading, adverse selection, electronic limit order book, short-term momentum, contrarian.

JEL classification: D82, D53, G01.

1 Toulouse School of Economics (CNRS-CRM).

This research was conducted within the FBF IDEI Chair on Investment Banking and Financial Markets. Bruno Biais benefited from the support of the European Research Council, grant 295484 TAP.

Introduction

In perfect markets, buyers and sellers immediately find each other and reap gains from trade at frictionless prices. Real markets, however, can fall short of delivering such welfare improvements, due to frictions.

Market frictions, indeed, can prevent final sellers from rapidly locating final buyers. In this context, intermediaries can provide liquidity to impatient sellers, by purchasing their assets and holding inventories, until they find final buyers. Such market-making services have been analysed theoretically by Ho and Stoll (1981, 1983), Grossman and Miller (1988) and Weill (2007).[2] What are the characteristics of intermediaries which enable them to supply liquidity? In fragmented markets, intermediation services can be provided by those agents with the best network linkages and the greatest search ability, which can be enhanced by high-frequency trading technology. Even if the market is centralised, delays can arise, reflecting that not all potential buyers and sellers are permanently monitoring the market, and that it takes time for investors to identify their trading needs, as analysed theoretically by Biais, Hombert and Weill (2014). In this context, as in that of fragmented markets, Gromb and Vayanos (2002) show that arbitrageurs, able to take positions in different markets, can provide valuable liquidity. Market-makers, however, bear costs when holding inventories, eg because they are risk–averse and reluctant to carry unbalanced inventory positions, as analysed theoretically by Ho and Stoll (1981), or because the principals of market-makers set position limits to discipline their agents. This suggests that the agents best placed to offer liquidity are those with the best inventory-holding ability, ie those with the greatest risk tolerance or the least acute agency problems. Because the aggregate inventory bearing capacity of market-makers is limited, however, liquidity shocks have a transient impact on prices, ie there are "limits to arbitrage," and liquidity supply is profitable (Shleifer and Vishny (1997), and Gromb and Vayanos (2002, 2010, 2015)).

Another market friction that restricts liquidity is adverse selection. As first shown by Akerlof (1970), adverse selection can magnify the price impact of trades and even lead to market breakdown. As shown by Glosten and Milgrom (1985) and Kyle (1985), adverse selection leads market-makers to post relatively high ask prices, and relatively low bid prices. Here again, however, the question arises: which agents will play the role of market-makers, and why? Efficiency suggests that the intermediaries should be the agents who are best able to mitigate adverse selection. Such ability could reflect better market-monitoring technology, enabling intermediaries to cancel their orders before they are picked off. This, however, could worsen the adverse selection problem for other investors, with less efficient monitoring technologies. Adverse selection for these investors could be further amplified if intermediaries took advantage of their timely market information to hit stale quotes themselves.

Since the beginning of the century, three developments have made these questions highly topical. First, equity markets have converged towards an electronic limit order book structure, in which a large number of different financial institutions (not just designated market-makers) can provide liquidity by leaving limit orders in the book. Second, low latency technologies have become available, increasing, at a cost, the ability to monitor changes in market conditions and react rapidly to them.

[2] Symmetrical arguments hold for the case of impatient buyers. In that case, market-makers can provide liquidity by selling assets they hold in inventory, or by short-selling the asset.

Third, regulatory reforms before the crisis contributed to the fragmentation of markets and the development of high–frequency trading, while regulatory reforms after the crisis made proprietary trading more costly and complex for investment banks.[3] How have these developments changed the economics of liquidity supply, and the gains from trades that can be reaped in financial markets? To shed light on these issues, we empirically analyse a new data set, with information about the orders and trades of different categories of members of Euronext, including proprietary traders and high-frequency traders.

Our data include a time-stamped record of all orders and trades (quantities and prices) on Euronext in French stocks during 2010. Our sample period brackets the Greek crisis of the summer of 2010, enabling us to analyse how liquidity supply compares between "normal" and crisis times. Our data also include anonymised member codes, and for each member, we know (i) the quality and speed of its connection to the market, and (ii) if its trades were 100% proprietary, 100% agency, or a mix of both. Using (i) we identify fast traders based on direct information about their technological investment, which contrasts with indirect identification, based on trading style. Because of its huge size, and also because of some technical characteristics of the Euronext market,[4] this data set is difficult to handle. At this stage, we have analysed 23 French stocks, including 10 large caps, nine mid-caps, and four small caps. The size of the corresponding data exceeds 7 tera–octets.

Our first main empirical finding relates to liquidity supply. We find that proprietary traders, whether fast or slow, tend to place marketable buy orders after price declines, and marketable sell orders after price increases. Thus, we find that proprietary traders follow contrarian strategies, buying against downward price pressure, and selling against upward price pressure. This contrasts with the other traders, who tend to buy after price rises and sell after price declines, which can be interpreted as momentum trading. While the proprietary traders' contrarian strategies rely on marketable orders, they supply liquidity to the market. This is consistent with proprietary traders being better able to carry inventory risk than other traders. This might be because they commit their own capital, rather than trading on other people's behalf. This could also reflect better incentive contracts. In both cases, superior ability to carry inventory risk would stem from having more "skin in the game."

Market liquidity could be hampered if liquidity supply evaporated in times of market stress (see eg Nagel 2012). Interestingly, the contrarian strategies of proprietary traders are particularly prevalent for small stocks, and during the Greek crisis. Thus, they are not "fair weather" liquidity suppliers, disappearing when the market most needs them. Moreover, the contrarian strategies of proprietary traders are on average profitable. This suggests they are able to identify when transient price pressure, possibly reflecting liquidity shocks, has driven prices away from equilibrium. In those circumstances, proprietary traders' marketable orders conduct risky

[3] Correspondingly, some investment banks had to move their prop trading activities to segregated subsidiaries, such as Descartes Trading for Société Générale.

[4] The majority of markets refresh the limit order book at the end of the trading day, eliminating unexecuted limit orders. This is not the case in Euronext, where a limit order can be left in the book for up to one year. Consequently there are many remaining limit orders, far from the quotes. This increases the size of the data set.

arbitrage. Thus, by absorbing selling or buying pressure, they tend to stabilise the market.

Our second major empirical finding relates to adverse selection. We find that the information content of marketable orders, although significantly positive, does not significantly differ across members' categories. In particular, fast traders' marketable orders do not seem to be more informed than those of slower traders. In contrast, we find that the adverse selection cost borne by non–immediately executed limit orders is significantly lower for fast proprietary traders. On the other hand, adverse selection costs are high for fast non–proprietary traders' limit orders. This suggests that technology, in itself, is not enough to mitigate adverse selection. It is also necessary that the traders have the incentives to use the technology efficiently. Moreover, while fast proprietary traders frequently trade with non–immediately executed limit orders, slow proprietary traders only infrequently do so. Thus, while slow proprietary traders mainly supply liquidity by placing contrarian marketable orders, fast proprietary traders also supply liquidity by placing non-immediately executed limit orders. We find, however, that this second type of liquidity supply becomes much less prevalent after the crisis.

Literature

Our paper relates to the rich literature empirically analysing algorithmic and high-frequency trading:

Hendershott, Jones and Menkveld (2011) show that liquidity provision in equity market is provided by algorithmic traders who are not officially designated as market-makers.

Brogaard, Hendershott and Riordan (2014) analyse NASDAQ data, relying on a classification of traders performed by the market organisers. They find that HFTs' marketable orders have information content, but not significantly more than other traders' marketable orders. They also find that HFTs' orders tend to be contrarian and are on average profitable. Thus, our empirical findings complement those of Brogaard, Hendershott and Riordan (2014) in two ways:

- Our findings relative to high-frequency traders are similar to theirs, and thus speak to the robustness of their results to a different market and a different way of identifying high-frequency traders.
- Our findings on proprietary traders are new, and thus provide an incremental contribution to the literature. We show that several characteristics of high-frequency traders' orders (ie contrarian, profitable and marketable) are also characteristics of the slow proprietary traders' orders in our data. This suggests that these traders are able to place contrarian and profitable orders not because they have low latency, but because they are proprietary traders.[5]

Brogaard, Hagstromer, Norden and Riordan (2014) find that the choice to co-locate is followed by reduced adverse selection costs. Our results on adverse selection are in line with their findings. Our incremental contribution, here, is to emphasise that

[5] The high-frequency traders identified by NASDAQ are very likely to also be proprietary traders.

technology may not be sufficient to reduce adverse selection. Appropriate incentives to use the technology efficiently seem to be also needed.

Hendershott and Menkveld (2014) offer a state space model to identify how price pressure relates to the specialist's inventory. In a sense, the proprietary traders in our data play a similar role to the specialist in Hendershott and Menkveld (2014).[6]

Our paper is also in line with analyses of liquidity supply by market participants that are not officially designated as market-makers. Franzoni and Plazzi (2014) document liquidity supply by hedge funds, classifying as liquidity supply purchases (or sales) at a price below (or above) a benchmark (eg the opening price or the VWAP). Nagel (2012) argues that the returns on short-term reversals can be interpreted as the returns from liquidity provision. Since, in our data, we observe the type of the market participants and the details of their orders, we can rely on more precise data to directly identify liquidity supply. Our empirical finding that proprietary traders supply liquidity with profitable contrarian trades confirms the relevance of the approach of Nagel (2012) and Franzoni and Plazzi (2014). Moreover, Nagel (2012) interprets the high return on liquidity supply strategies when the VIX is high as evidence that liquidity supply is reduced at those times. Our direct evidence on liquidity supply offer some nuance on this point: on the one hand, we find that proprietary traders continued to place profitable marketable contrarian orders during the Greek crisis of the summer of 2010, so, in that sense, liquidity did not "evaporate". On the other hand, we find that the crisis triggered a drop in the placement of non-immediately executed limit orders in small caps by fast proprietary traders. Other things equal, this corresponds to a decline in liquidity supply.

The next section presents our data and summary statistics. Section 3 analyses the placement and informational content of marketable orders. Section 4 analyses the adverse selection costs incurred by non-immediately executed limit orders. Section 5 summarises our results and discusses policy implications.

Market structure, data and summary statistics

Euronext is an electronic limit order market, operating continuously for most of the stocks. Orders are submitted by the financial institutions that are members of Euronext. Members are assigned an identification code. Most institutions have a single ID code, but some choose to have up to four different codes to differentiate their activities (eg brokerage versus proprietary trading). Each ID code can have several links to the exchange. The total fee charged by Euronext to a member depends on the number of links connected to the Euronext servers, and on the capacity of each link. This capacity is named "throughput" and indicates the maximum number of messages (order submission, cancellation or modification of an existing order) that the traders are allowed to send to the Exchange in one calendar second via this link. The throughput of each link can be 10, 25, 50, 100, 200, 250, or 500. It seems that a large throughput does not only increase the number of messages that can be sent, but also increases the speed with which they reach the market. We are able to document this empirically because, in our data, we observe (at the microsecond level)

[6] Furthermore, Kirilenko, Kyle, Samadi and Tuzun (2014) find that HFTs did not cause the Flash Crash but exacerbated it. Menkveld (2013) document liquidity supply by large HFTs across ChiX & Euronext.

the time at which an order was sent to the market and the time at which the order was executed. We also observe the throughput of the link via which the order was sent. Marketable orders, by construction, are immediately executed. Thus, for these orders, the difference between the execution time and the submission time is an inverse measure of the speed of the connection. In Graph 1, panel A, we plot for each link the median of this time difference against the throughput of the link. The Graph illustrates that links with higher throughput tend to also have a higher connection speed. Moreover, Euronext further proposes to its members to co-locate their servers close to its own.[7] Co-location fees depend on the physical space used by the member, and on the electricity consumption of the servers.

We obtained our data from the Autorité des Marchés Financiers (the French financial markets regulator) and Euronext. So far, we have analysed 23 French stocks between February 2010 and August 2010. The sample includes 10 large caps (one financial and nine non-financial with float between 1,048 and 3,884 million euros), nine mid-caps (one financial and eight non-financial, with float between 181 and 960 million euros), and four small caps (all non-financial, with float between 51 and 145 million euros).

Some 150 Euronext member IDs traded these stocks during the sample period. Some 28 of them only traded on their own account, and we classify them as prop traders. Thirty-seven traded only on behalf of customers and we classify them as pure brokers. Eighty-five conducted some trades on their own account and some trades on behalf of customers. We classify them as dual traders.[8]

As mentioned above, we observe the number of links for each member ID, and the throughput of each link. For each member ID, we compute the messaging capacity, ie the sum of the throughputs of all the links. As discussed above, in line with Graph 1, Panel A, messaging capacity is also correlated with speed of connection. Graph 1, Panel B, illustrates the distribution across member IDs of messaging capacities. They range between 0 and 4600. Traders with higher capacity and higher connection speed can observe changes in market conditions faster and react more quickly to them. We chose to classify as fast traders the 17 members with capacity above 1,300 messages per second. We checked that our qualitative results were robust to changing the threshold, eg setting it at 800 or 1,600.

Out of the 17 fast traders, six are proprietary traders while the 11 others are dual traders. Members' ID codes are anonymised, so that we do not observe the identity of the traders, nor do we have direct information about their type. We do, however, have some indirect information, unrelated to the data set, based on which we formed conjectures. It is highly likely that the six fast proprietary traders are high-frequency trading "boutiques", similar to the high-frequency traders identified in the Nasdaq data set. Moreover, it is likely that the 11 fast dual traders are, typically, European banks, sending proprietary trades as well as agency trades to the exchange via the same membership channel.

7 In May 2010, Euronext servers were moved to Basildon (in the neighbourhood of London). Members can choose to co-locate some, all or none of their links.

8 Actually, for each trade we observe if it was reported by the member as proprietary or agency. But we have been told by market participants that these reports can be quite noisy. To reduce the impact of such noise, we chose to classify as pure proprietary traders those reporting only proprietary trades and as pure brokers those reporting only agency trades.

Out of the 133 slow traders, 22 are proprietary traders, 74 are dual traders, and 37 are pure brokers. Once again, because the data set is anonymised, we have no direct information about who these members are. It is likely, however, that the 22 slow proprietary traders include proprietary trading desks of large investment banks with their own membership and possibly some hedge funds.

Combining the two criteria, we classify the members in our sample into five categories: fast prop–traders, slow prop–traders, fast dual traders, slow dual traders, and slow brokers. Note that no pure broker in our sample was fast. This does not mean they had absolutely no high–frequency trading technology. It is likely they also rely on algorithms, eg to search for best execution. Their connections to the market, however, are less advanced than those of the fast traders.

As mentioned above, our sample period brackets the Greek crisis of mid-2010. Graph 2 plots the evolution of the VIX volatility index during our sample period. The Vix jumps on 23 April, when Greece asks for a bailout, and remains elevated until late June. Thus, we split our sample in three subperiods: The first period, before the crisis, is from 23 February to 22 April. The second period, corresponding to the crisis, is from 23 April to 22 June. The third period, after the crisis, is from 23 June to 23 August.

Graph 3 depicts the number of trades per member, per stock, per day. Graph 3, Panel A, shows that fast traders (both proprietary and dual) trade more often, and rely more on non-immediately executed limit orders than do slow traders. Slow proprietary traders trade less than fast traders, but more than other slow traders do. Moreover, unlike fast proprietary traders, they rely mainly on marketable orders.

Panels B and C of Graph 3 show how these results vary with market capitalisation and periods. Comparing Panel B (large caps) and Panel C (small and mid-caps), the number of trades is much larger for the former than for the larger. Furthermore, for large caps the number of trades increases during the crisis, but the relative frequencies of marketable orders and non-immediately executed limit orders are rather stable through time. For small caps, however, the behaviour of fast proprietary traders is quite different. First, their trading volume does not increase during the crisis. Second, while before the crisis they frequently traded via non-immediately executable limit orders, during and after the crisis they considerably reduce their reliance on this type of trade. This is an important observation, to which we come back below, when we analyse how the crisis affected adverse selection costs.

Graph 4 offers a graphical illustration of the frequency with which each category of traders' marketable orders hit each category of traders' limit orders. Each colour corresponds to the category of the trader whose marketable order was executed. For example, the plain red bar corresponds to the marketable orders placed by fast proprietary traders. It shows that, when a fast prop trader places a marketable order, 19% of the time it hits another fast trader, 8% of the time it hits a slow prop trader, 41% of the time it hits a fast dual trader, 27% of the time it hits a slow dual trader, and 4% of the time it hits a pure broker's order. Interestingly, the frequencies are very similar for marketable orders placed by other categories of traders. So, it is not the case that fast traders "target" a certain category of trader's limit orders.

Marketable orders

To estimate the information content of orders in the simplest possible way, we use the percentage change from the midquote just before the trade to the midquote 2 minutes after the trade. Thus, for a trade taking place at time t, the information content is

$$\frac{M_{t+2} - M_{t-}}{M_{t-}} * sign_t ,$$

Where M_{t^-} denotes the midquote just before the trade, M_{t+2} denotes the midquote 2 minutes after the trade, and $sign_t$ takes the value 1 if the time–t marketable order is a buy order, and –1 if the time–t marketable order is a sell order.

Graph 5, Panel A, depicts the informational content of the marketable orders placed by different categories of traders. After a marketable buy (or sell) order placed by a fast proprietary trader, the midquote increases (or decreases) on average by 3.5 basis points. The informational content of the marketable orders placed by other categories of traders is not very different, ranging between 3.4 and 4.1 basis points. This suggests that the marketable orders placed by fast traders are not more informed than the marketable orders placed by other traders. Correspondingly, they do not generate more adverse selection for limit orders standing in the book.

Graph 5, Panel B, depicts how the informational content of marketable orders varies across stocks and through time.[9] It shows that, for all traders, the informational content of orders is larger for small caps and during the crisis, and remains higher after the crisis than before. Also, while, as shown in Graph 5, Panel A, fast traders' orders are, in general, not better informed than other traders' orders, during the crisis they are.

Graph 6 depicts the evolution of the midquote five minutes before and after the placement of a marketable order by the different categories of traders, ie it plots the average of

$$\frac{M_{t+h} - M_{t-}}{M_{t-}} * sign_t ,$$

where h varies from –5 minutes to +5 minutes. Panel A shows the results for large caps, while Panel B shows the results for small caps. The two panels confirm the above discussed findings that, after the trade, the information content of marketable orders is similar for all categories of members, and is higher for small stocks. The new information in Graph 6, relative to Graph 5, is about what happens before the trade. Graph 6 shows that dual traders or brokers place marketable buy (or sell) orders after price increases (or decreases). This is consistent with dual traders and brokers riding short–term momentum waves, or splitting orders. In stark contrast, fast proprietary traders (be they fast or slow), buy after price decreases and sell after price increases.

9 The information content of fast prop traders' orders is larger than that of other traders for all subcases in Panel B. This does not contradict the fact that it is not higher in Panel A, because fast proprietary traders are relative more active in large cap (for which the informational content of orders is lower) than in small caps.

That is, they follow contrarian strategies. Thus their marketable orders supply liquidity to the market, helping it accommodate buying or selling pressure.

Is such liquidity supply robust ? Or does this liquidity evaporate when the market really needs it, eg in times of crisis? To shed light on this, Graph 7 depicts the evolution of the midquote five minutes before and after the placement of a marketable order, similarly to Graph 6, but distinguishing between the period before the crisis (Panel A), during the crisis (Panel B) and after the crisis (Panel C). The Graph shows that there are no qualitative differences across periods. During the three periods, dual traders and pure brokers buy after price rises and sell after price drops. And, during the three periods, proprietary traders buy after price drops and sell after price rises. The only difference is that, during the crisis, the magnitude of the price changes is larger. Thus, the liquidity supplied to the market by the contrarian orders of proprietary traders does not seem to evaporate during times of market stress.

Our results are consistent with those of Brogaard, Hendershott and Riordan (2014), who find that HFTs' marketable orders tend to be contrarian. Indeed, the HFT firms identified by Brogaard, Hendershott and Riordan (2014) are likely to be proprietary traders. The new finding in the present study is that slow traders also place such contrarian marketable orders, when they are proprietary traders. This suggests that what enables traders to conduct such contrarian strategies is maybe not technology but the ability to trade on one's own account. That ability reduces agency conflicts between investors and traders, and thus increases the ability to carry inventory, and, correspondingly, supply liquidity to the market.

While such liquidity supply could be beneficial to the market, by accommodating liquidity shocks, one might wonder whether it is sustainable. In particular, is it profitable? To shed light on this, we computed the average profits of the marketable orders placed by the different categories of market participants. To estimate the profitability of orders in the simplest possible way, we use the percentage difference between the transaction price and the midquote two minutes after the trade. Thus, for a marketable order executed at time t, the profit is

$$\frac{M_{t+2} - P_t}{M_{t-}} * sign_t ,$$

where P_t denotes the transaction price, M_{t+2} denotes the midquote two minutes after the trade, and $sign_t$ takes the value 1 if the time–t marketable order is a buy order, and –1 if the time–t marketable order is a sell order.

Graph 8 depicts the results. Panel A of Graph 8 shows that proprietary trader's marketable orders earn positive profits: 0.8 bp on average for fast proprietary traders, and 1.7 bp on average for slow proprietary traders. This suggests that the ability of proprietary traders to supply liquidity with contrarian orders is sustainable, in the sense that it is profitable.[10]

Panel A of Graph 8 shows that, in contrast, the momentum-riding strategies of other traders appear to be unprofitable. In particular, slow dual traders lose on average around 1.1 bp per trade, and slow brokers almost 4 bp. Thus, while these orders trade in the direction of market movement, buying before price increases and

[10] Proprietary trades are exempted from trading fees on Euronext, so the above-mentioned profitability is not eliminated by exchange fees.

selling before declines, they are not profitable, because the spread they have to pay exceeds their informational content.

Panel B of Graph 8 shows how the profitability of marketable orders vary with the market capitalisation of the traded stocks and with the crisis. Both the profits of the proprietary traders and the losses of the other traders are larger for small caps than for large caps. Furthermore, the profitability of fast proprietary traders' marketable orders increase during the crisis, while those of slow proprietary traders decrease somewhat for small caps, but remain largely positive. This is consistent with the results depicted in Graph 7: proprietary traders continue to supply liquidity to the market via contrarian marketable orders during the crisis, and this activity continues to be profitable.

As mentioned above, the behaviour we interpret as momentum riding could in fact reflect order-splitting. We can address this issue because our data contain the ID codes of market participants (anonymised of course.) If order-splitting was the reason why dual traders and slow brokers bought after price rises (and sell after drops), we should see that, after a dual trader or broker placed a marketable order, the next marketable order would often be in the same direction and stemming from the same trader. To offer evidence on this point, we estimate a contingency table in the same spirit as in Biais, Hillion and Spatt (1995). Graph 9 presents the probability that, after a marketable order was placed by a given category of participant, the next marketable order is placed by the same participant or another one, and is in the same direction or the opposite one. The Graph shows that, after a marketable order was placed by a fast dual trader, the probability that the next marketable order stems from the same trader is below 10%. For slow dual traders, this probability is between 10% and 15%, for slow brokers it is between 15% and 20%. These frequencies are not very large, suggesting that order-splitting is not the main driving force, and momentum trading is likely to be at play. Interestingly, after a marketable order was placed by a fast prop trader, the probability that the next marketable order stems from the same trader is very low (around 5%), while after a marketable order from a slow prop trader it is very large (above 35%). This suggests, that, at least in the case of slow prop traders, order-splitting and contrarian liquidity supply coexist. Another implication from Graph 9 is that order-splitting seems to be more prevalent for slow traders than for fast traders.

The overall message emerging from the above findings is the following: slow traders tend to split orders, while fast traders engage less in order-splitting. Proprietary traders help the market absorb liquidity shocks by placing contrarian marketable orders, while other traders tend to consume liquidity by riding momentum waves.

Adverse selection costs incurred by limit orders

Limit orders left in the book are exposed to adverse selection, as analysed, eg by Glosten and Milgrom (1985), Glosten (1994) and Biais, Martimort and Rochet (2000). Fast trading technology, however, enhances traders' ability to monitor market movements. This may enable them to cancel or modify stale quotes, before they receive adverse execution.

To estimate the adverse selection cost incurred by limit orders in the book, we use the percentage change from the midquote just before the trade to the midquote two minutes after the trade. Thus, for a trade taking place at time t, the information content is

$$\frac{M_{t+2} - M_{t-}}{M_{t-}} * sign_t^{\text{limit}}$$

where M_{t^-} denotes the midquote just before the trade, M_{t+2} denotes the midquote 2 minutes after the trade, and $sign_t^{\text{limit}}$ takes the value 1 if the limit order hit at time-t is a buy order, and –1 if it is a sell order. This is similar to the measure of information content illustrated in Graph 5, but, while in Graph 5 we condition on the type and direction of the aggressive marketable order, in Graph 10 we condition on the type and direction of the passive limit order (hit by a marketable order).

Graph 10, Panel A, shows that, after a limit buy (or sell) order left in the book by a fast proprietary trader is executed, the midquote decreases (or increases) on average by 2.8 basis points. The adverse selection costs of limit orders left in the book by other categories of traders are 3.1 bp for slow proprietary traders, 3.8 bp for fast dual traders, 4.1 bp for slow dual traders and 5 bp for slow brokers. Thus, the adverse selection cost is lowest for fast proprietary traders. Adverse selection costs are higher for fast non-proprietary traders. This suggests that technology, in itself, is not enough to mitigate adverse selection. It is also necessary that the traders have the incentives to use the technology efficiently.

Graph 10, Panel B, compares adverse selection costs before, during and after the crisis, and also across large and small caps. Overall, adverse selection costs are higher for small caps. They are also higher during the crisis. There are some differences between the different categories of traders:

- For non-proprietary traders, adverse selection costs increase during the crisis, and remain elevated afterwards.
- For slow proprietary traders, there is a similar pattern, except that adverse selection costs decline after the crisis.
- For fast proprietary traders, while the pattern is similar for small and mid-caps, for large caps adverse selection costs remain low throughout the period, and are not higher during or after the crisis. This is consistent with the patterns depicted in Panels B and C of Graph 3: fast proprietary traders seem to be able to control adverse selection costs for large caps during the crisis, and thus continue to supply liquidity in these stocks via limit orders. In contrast, for small and mid-caps, their adverse selection costs increase and they reduce their limit order liquidity supply. This is a form of "liquidity evaporation".

Graph 11 depicts the profits earned by limit orders left in the book, estimated as

$$\frac{M_{t+2} - P_t}{M_{t-}} * sign_t^{\text{limit}},$$

where P_t denotes the transaction price, M_{t+2} denotes the midquote two minutes after the trade, and $sign_t^{\text{limit}}$ takes the value 1 if the limit order hit at time-t is a buy order, and –1 if it is a sell order. Panel A of Graph 11 shows that the limit orders of fast proprietary traders (and also those of slow dual traders) earn slightly positive profits on average. That is, for these orders, the bid–ask spread (which they earn) is slightly larger than the adverse selection cost (which they incur). In contrast, the limit orders left in the book by dual fast traders lose 3 bp per trade on average, and those of slow proprietary traders 8 bp per trade. That is, for these limit orders the adverse selection cost is so large that it exceeds the spread. This is consistent with (i) our observation above that slow proprietary traders rarely use non–immediately executable orders and (ii) the conjecture that they are aware that such orders would be loss-making for them.

Panel B of Graph 11 shows that the profitability of non-immediately executed limit orders varies considerably with capitalisation and period. Both losses and profits are much lower for large caps. The most striking result is the large profitability of fast proprietary traders' limit orders in small caps during the crisis and afterwards. These are the most difficult stocks and times, for which adverse selection costs are the largest, as shown in Graph 10, Panel B. Yet, fast proprietary traders apparently are able to earn the spread to such an extent that they more than offset these costs.

The results in Graphs 10 and 11 suggest that fast proprietary traders' limit orders are less adversely selected than others, in line with the notion that fast proprietary traders monitor the market and often cancel or modify their orders before receiving adverse execution. To document this point, we computed the number of cancellations and the number of updates to less aggressive quotes for each category of trader, normalising these numbers by the number of trades. Graph 12 shows that, for fast proprietary traders, the number of updates to less aggressive quotes is 3.7 times as large as the number of trades outside the crisis, and 4.2 times during the crisis. For cancellations, these numbers are 12.8 outside the crisis, and 15.1 during the crisis. These numbers are much larger than the corresponding numbers for fast dual traders and much, much larger than the corresponding numbers for other categories of traders. That the frequency of cancellations and modifications by fast proprietary traders is higher during the crisis suggests that these traders made greater efforts to monitor the market and adjust to it. The cost of this effort increased the cost of leaving limit orders in the book during the crisis.

Conclusion and policy implications

We analyse a unique data set enabling us to observe, in addition to time-stamped orders and trades on Euronext, the number of messages per second traders can exchange with the market, and whether they are prop traders. Our sample period in 2010 brackets the Greek crisis in mid-2010.

We find that proprietary traders, whether fast or slow, tend to place marketable buy orders after price drops, and marketable sell orders after price increases. And, after they have bought, the market tends to recover while, after they have sold, the market tends to retreat. This is consistent with proprietary traders helping the market to accommodate liquidity shocks, and thus reducing transient deviation from efficient pricing. Moreover, we find that this liquidity supply remains available during the crisis. In that sense, liquidity supply by proprietary traders does not evaporate when it is most needed.

We also find that the limit orders of fast proprietary traders (but not other fast traders) incur lower adverse selection costs than the limit orders of other traders. Our results suggest that fast proprietary traders supply liquidity to the market via contrarian marketable orders and non-immediately executed limit orders, while slow proprietary traders supply liquidity only via contrarian marketable orders. The crisis, however, triggered a decline in the placement of non-immediately executable orders for small and mid-caps by fast proprietary traders. In that sense, some liquidity "evaporated".

Our empirical findings suggest that current regulatory reforms might have unintended negative consequences:

- Under MIFID 2, trading venues will be required to cap the ratio of the number of messages to the number of trades by any given participant. This might be counterproductive, as our findings suggest that fast proprietary traders rely on numerous cancellations and updates to reduce the adverse selection cost incurred by their limit orders. Capping the percentage of cancellations and updates could increase the adverse selection costs incurred by limit orders left in the book, and thus deter the provision of liquidity by these orders. This could be harmful for market liquidity, especially at times of market stress, when the need to modify and cancel orders is particularly acute.
- In this context, new banking regulations, making it more difficult and costly for banks to engage in proprietary trading, might also reduce market liquidity.

References

Biais, B, J Hombert and P Weill (2014): "Equilibrium pricing and trading volume under preference uncertainty", *Review of Economic Studies*, vol 18, no 4.

Biais, B, P Hillion and C Spatt (1995): "An empirical analysis of the limit order book and the order flow in the Paris Bourse", *Journal of Finance*, vol 50, no 5, pp 1655–89.

Biais, B, D Martimort and J Rochet (2000): "Competing mechanisms in a common value environment", *Econometrica*, vol 68, pp 799–838.

Brogaard, J, T Hendershott and R Riordan (2014): "High-frequency trading and price discovery", *Review of Financial Studies*, vol 27, pp 2267–306.

Brogaard, J, B Hagströmer, L Nordén and R Riordan (2014): "Trading fast and slow: colocation and market quality", working paper, University of Stockholm.

Franzoni, F and A Plazzi (2014): "What constrains liquidity provision? Evidence from hedge fund trades", working paper, Swiss Finance Institute.

Glosten, L (1994): "Is the electronic open limit order book inevitable?", *Journal of Finance*, vol 49, pp 1127–61.

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

Gromb, D and D Vayanos (2010): "Limits of arbitrage: the state of the theory", *NBER Working Paper*, no 15821.

Gromb, D and D Vayanos (2002): "Equilibrium and welfare in markets with financially constrained arbitrageurs", *Journal of Financial Economics*, vol 66, pp 361–407.

——— (2015): "The dynamics of financially constrained arbitrage", working paper, London School of Economics.

Grossman, S and M Miller (1987): "Liquidity and market structure", *Journal of Finance*, pp 617–33.

Hendershott, T, C Jones and A Menkveld (2011): "Does algorithmic trading improve liquidity?", *Journal of Finance*, vol 66, pp 1–33.

Ho, T and H Stoll (1983): "The dynamics of dealer markets under competition", *Journal of Finance*, vol 38, pp 1053–74.

——— (1981): "Optimal dealer pricing under transactions and return uncertainty", *Journal of Financial Economics*, vol 9, pp 47–73.

Nagel, S (2012, "Evaporating liquidity", *Review of Financial Studies*, vol 25, pp 2005–39.

Shleifer, A and R Vishny (1997): "The limits of arbitrage", *Journal of Finance*, vol 52, pp 37–55.

Weill, P (2007): "Leaning against the wind", *Review of Economic Studies*, vol 74, pp 1329–54.

Panel A: Connection speed plotted against throughput

Graph 1

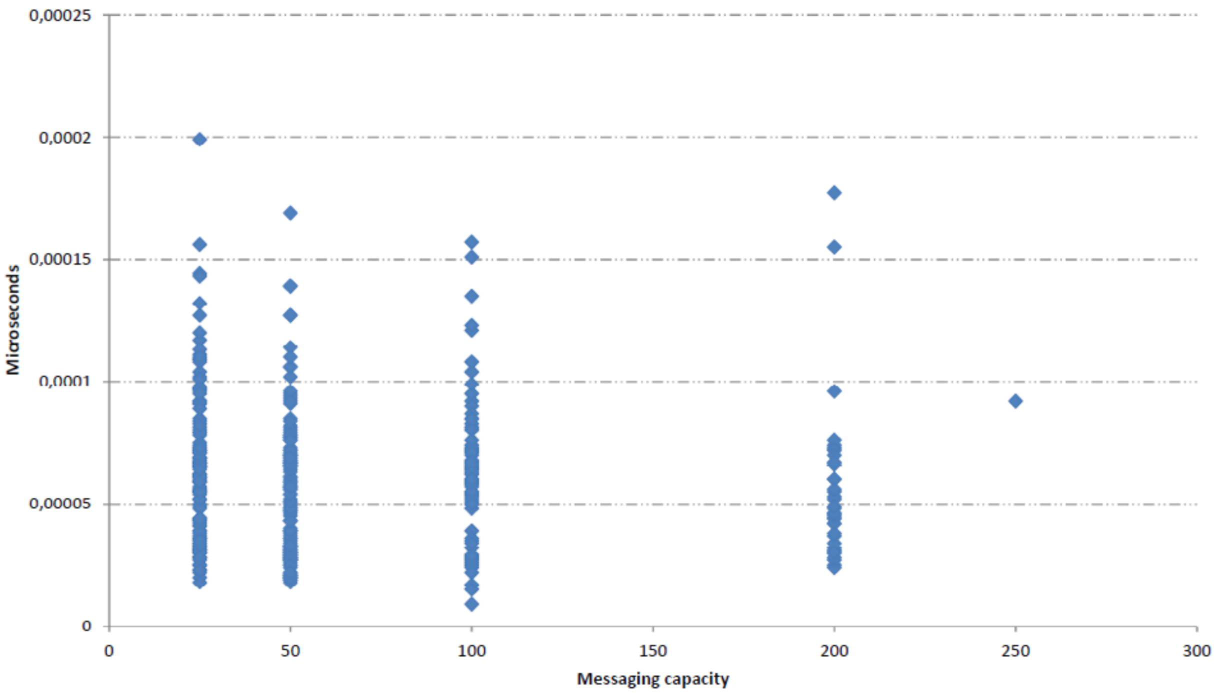

Panel B: Distribution of trading capacity across members

Graph 1

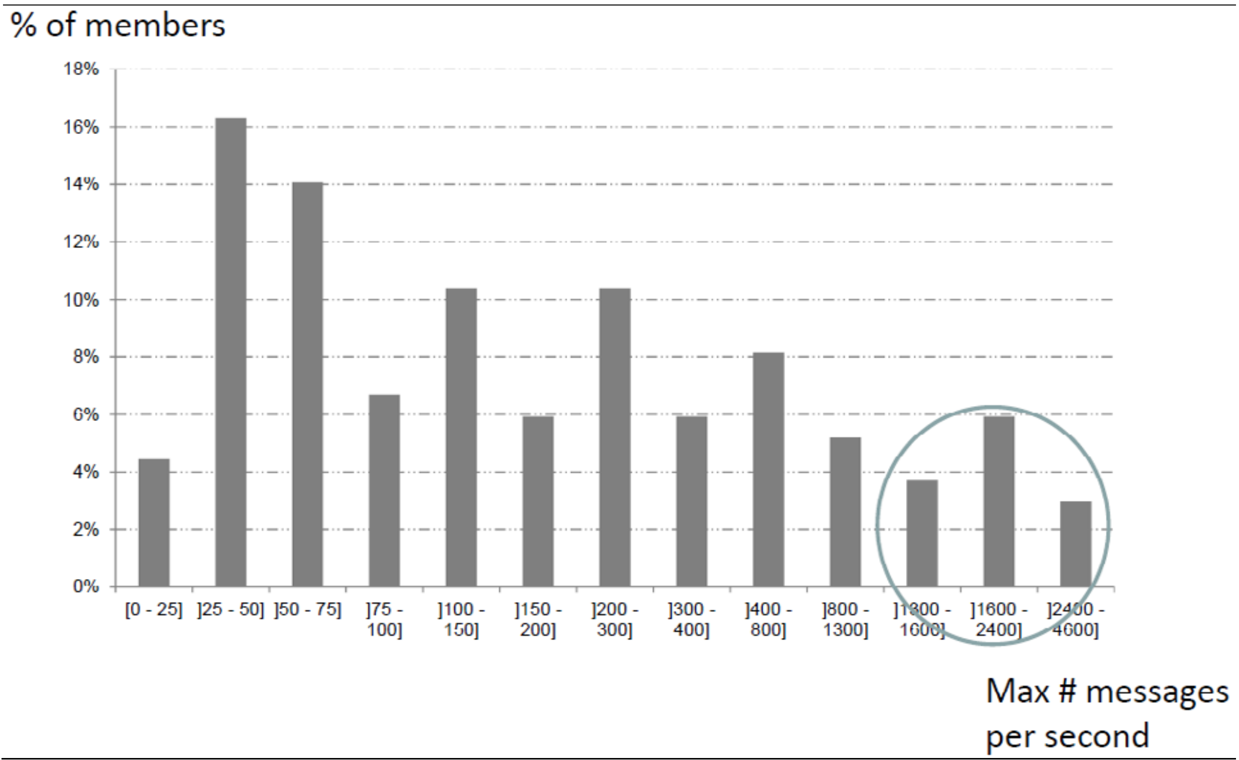

VIX during the sample period | Graph 2

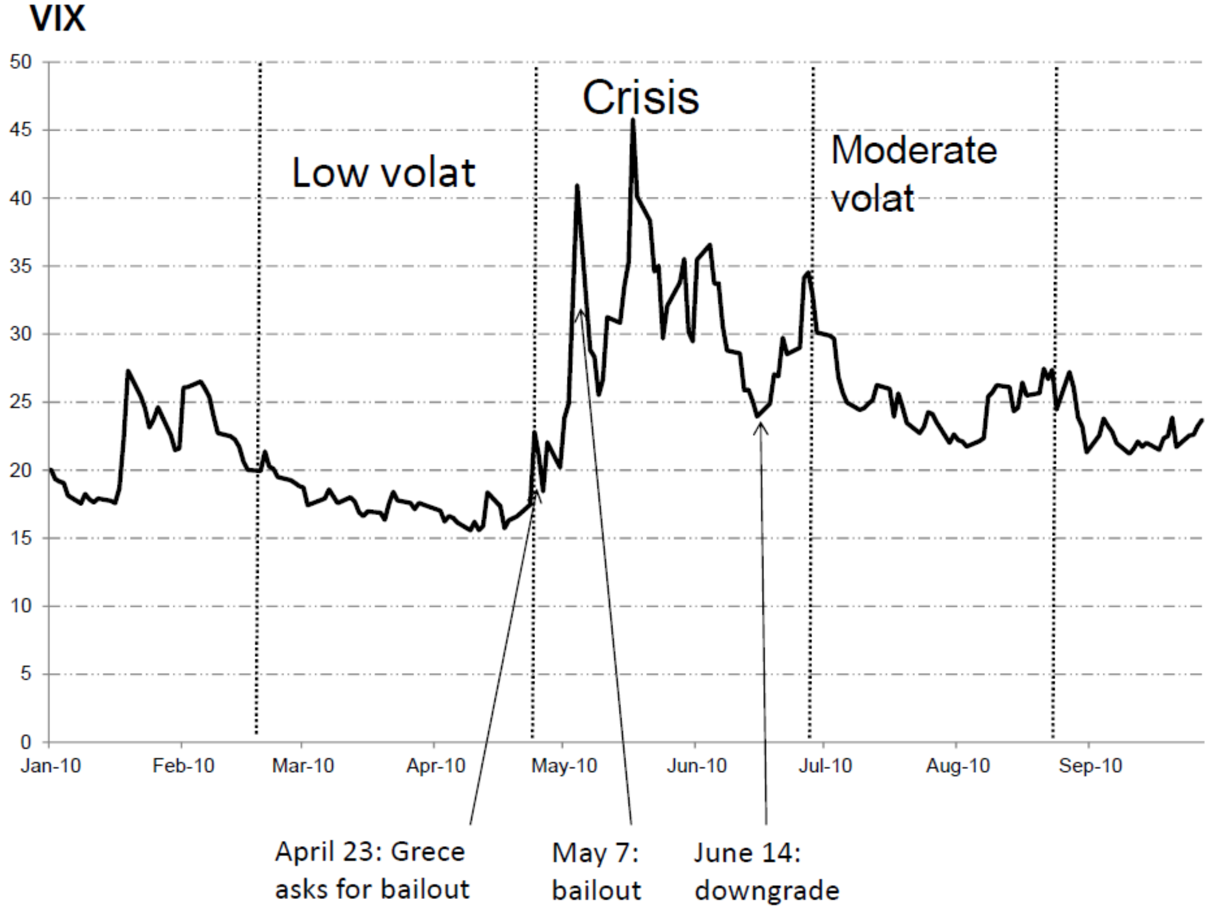

Panel A: Number of trades per member, stock and day | Graph 3

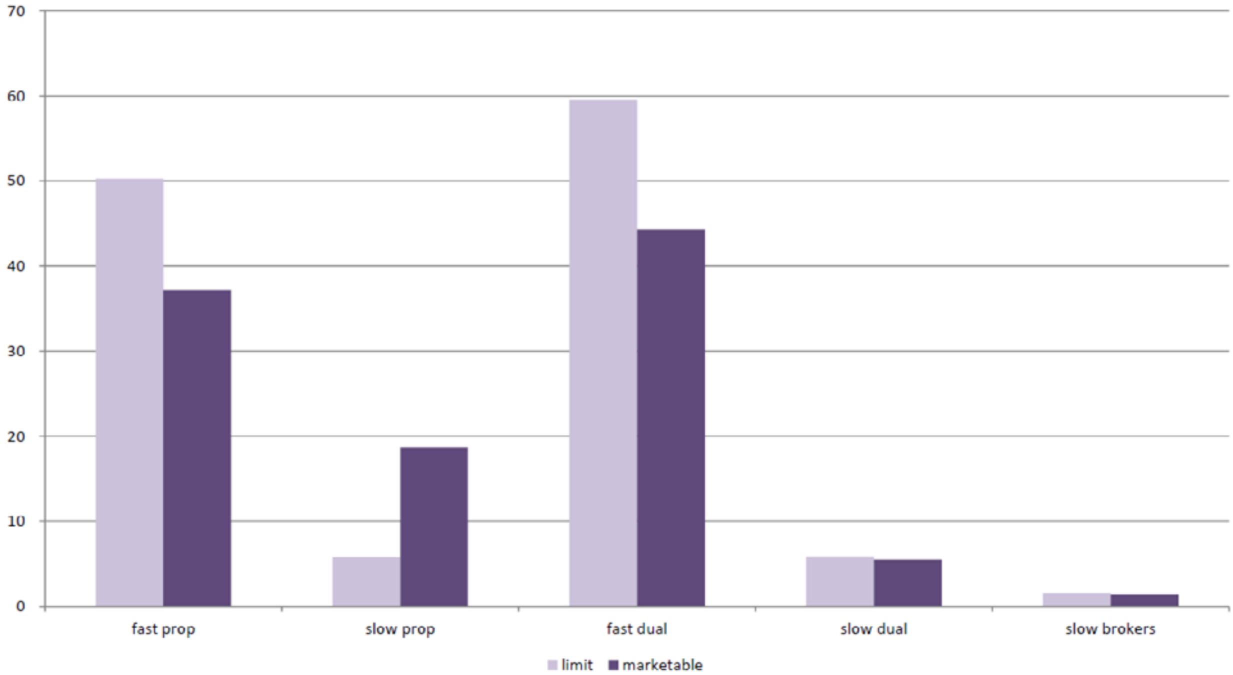

Panel B: Number of trades per member, stock and day – Large caps by period

Graph 3

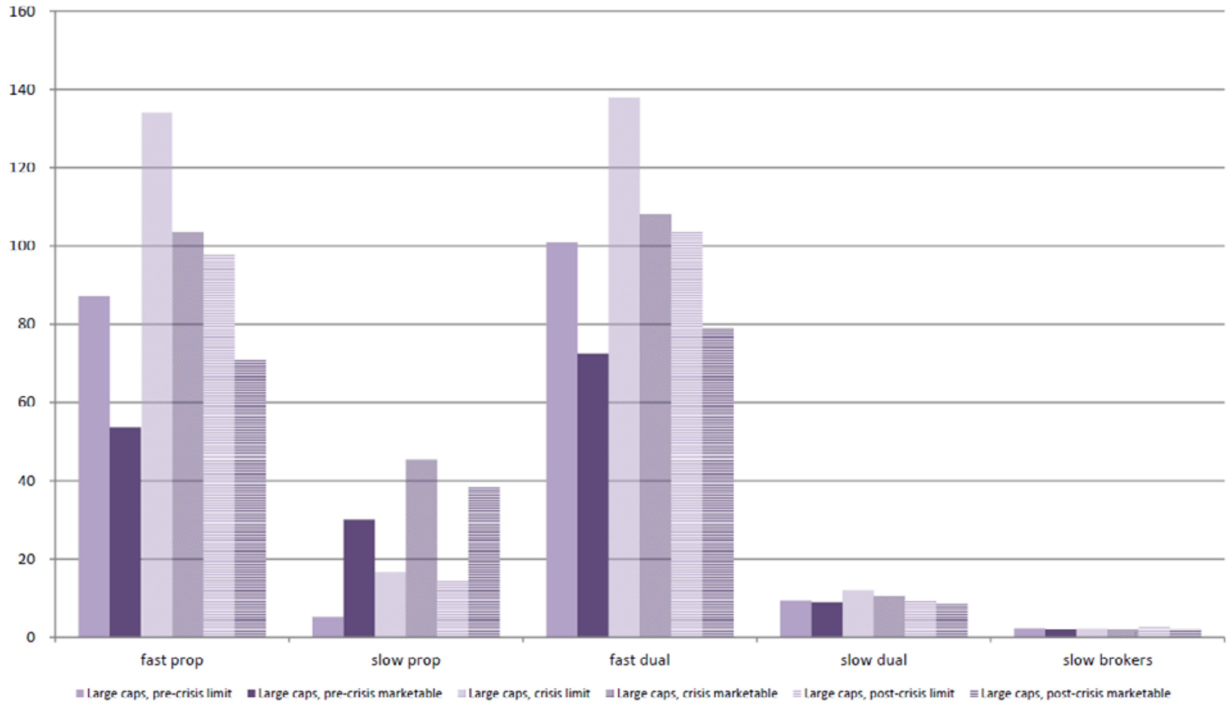

Panel C: Number of trades per member, stock and day – Small and mid-caps by period

Graph 3

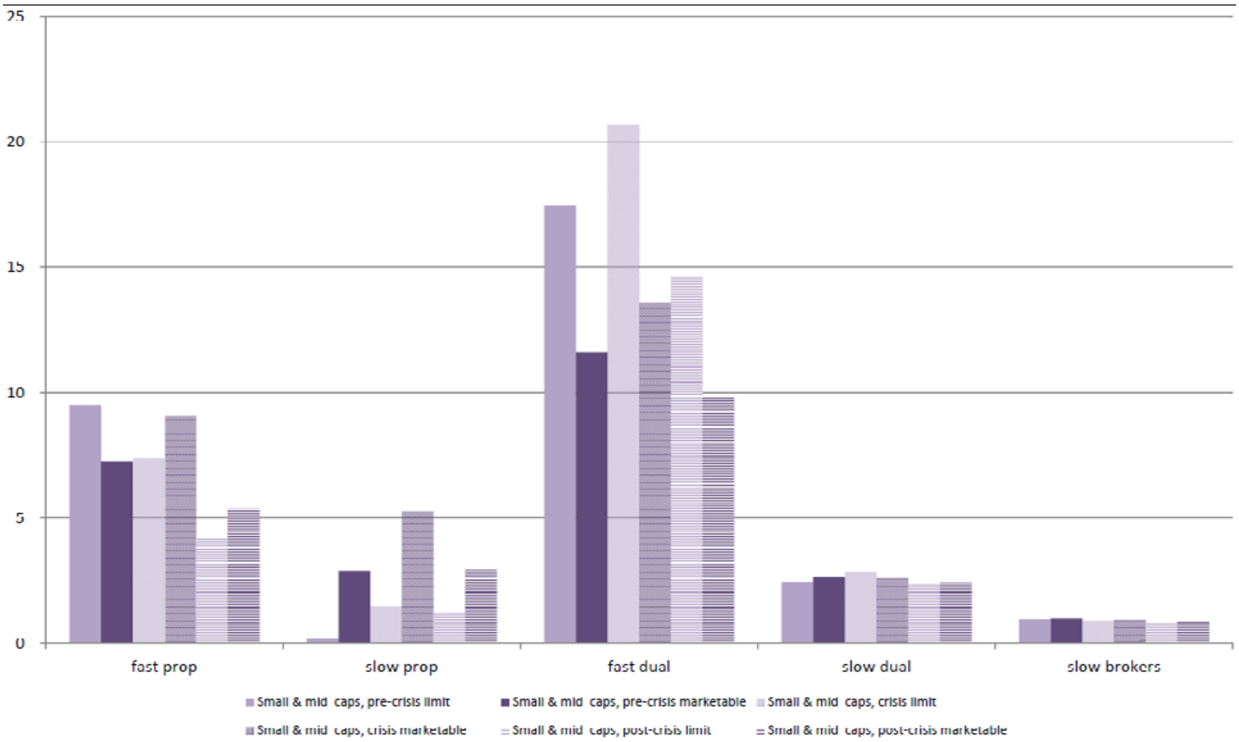

Who trades with whom? Graph 4

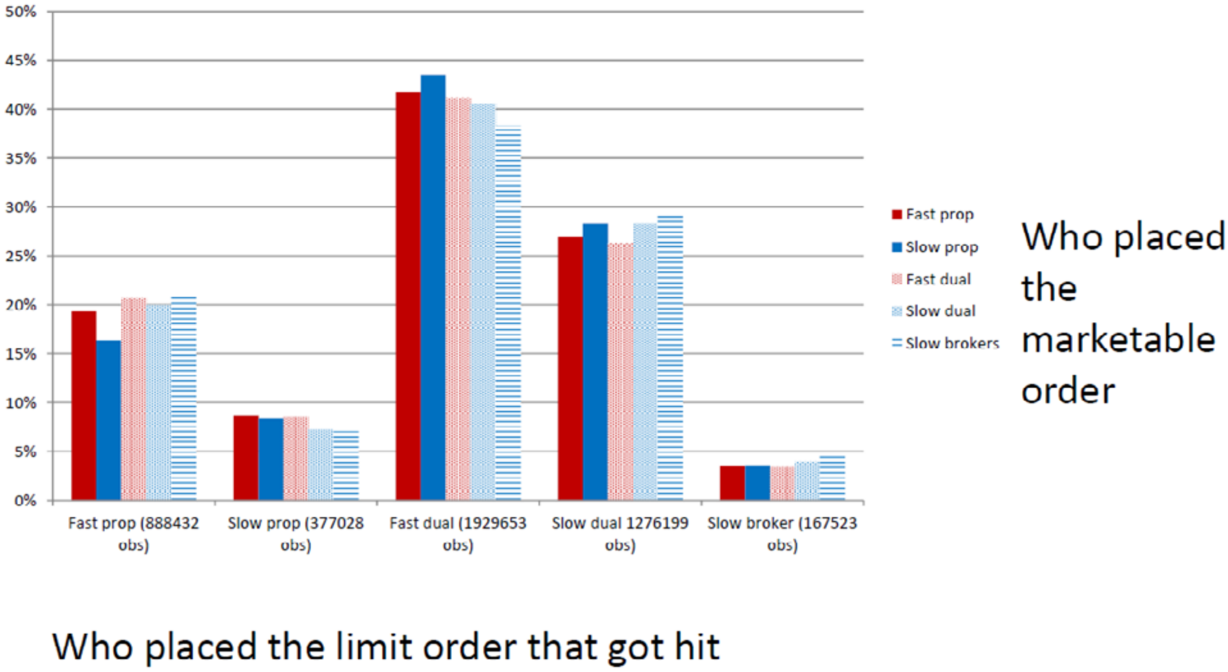

Panel A: Average information content of marketable orders

$(M_{t+2min} - M_{t-})/(M_{t-})$ * (sign of marketable order) * 100 Graph 5

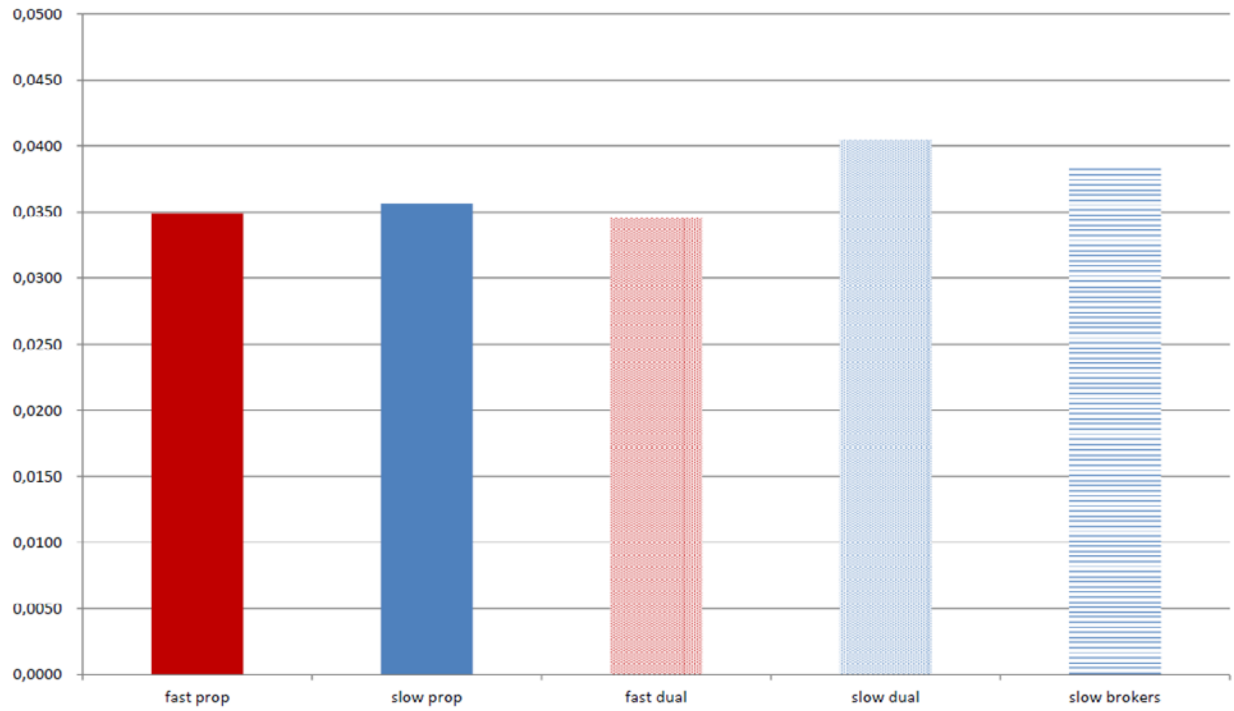

Panel B: How does the information content of marketable orders vary?

$(M_{t+2min} - M_{t-})/(M_{t-})$ * (sign of marketable order) * 100 Graph 5

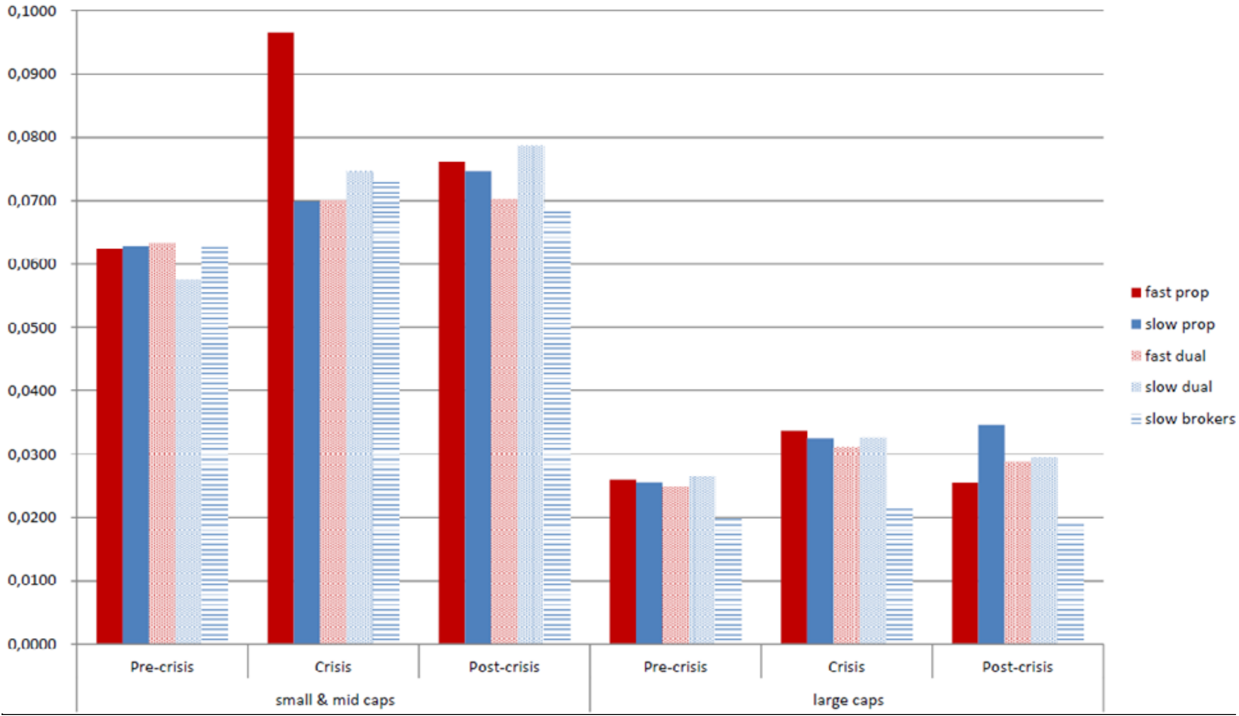

The midquote around marketable order executions Graph 6

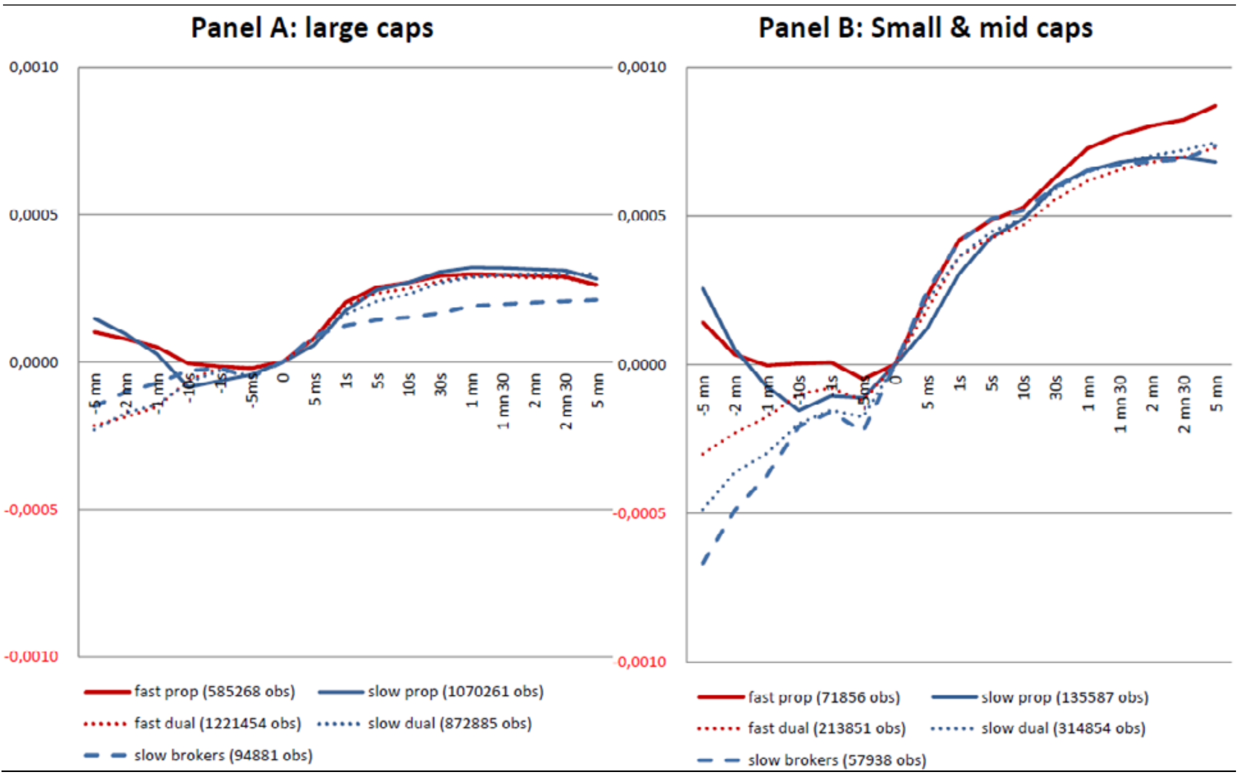

Momentum and contrarian marketable orders during crisis Graph 7

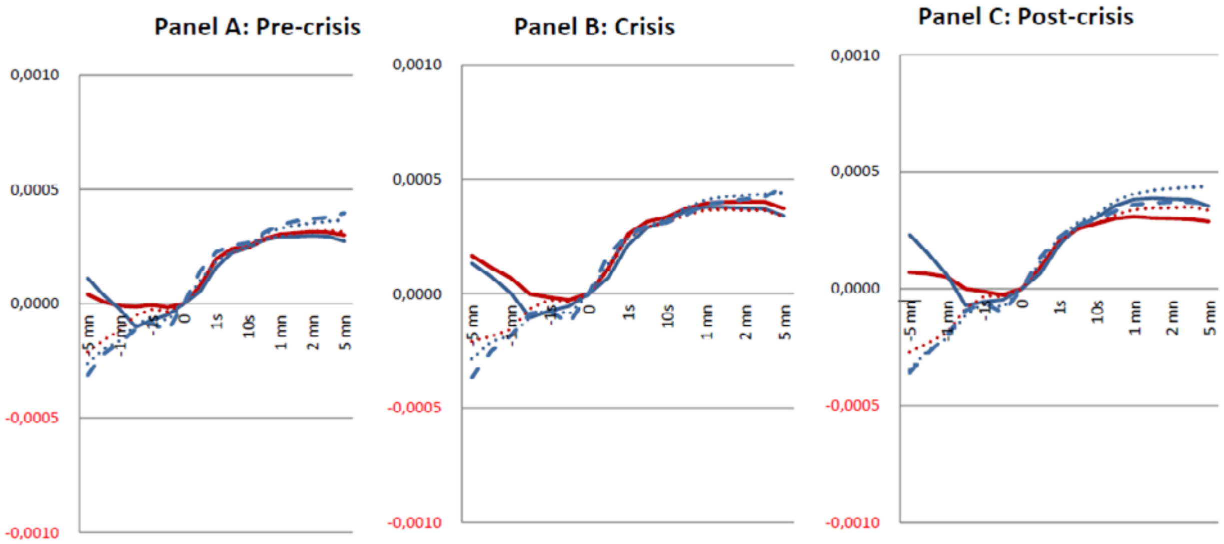

Panel A: Marketable orders' profits

(Mt+2min – Pt)/(Mt-) * (sign of marketable order) * 100 Graph 8

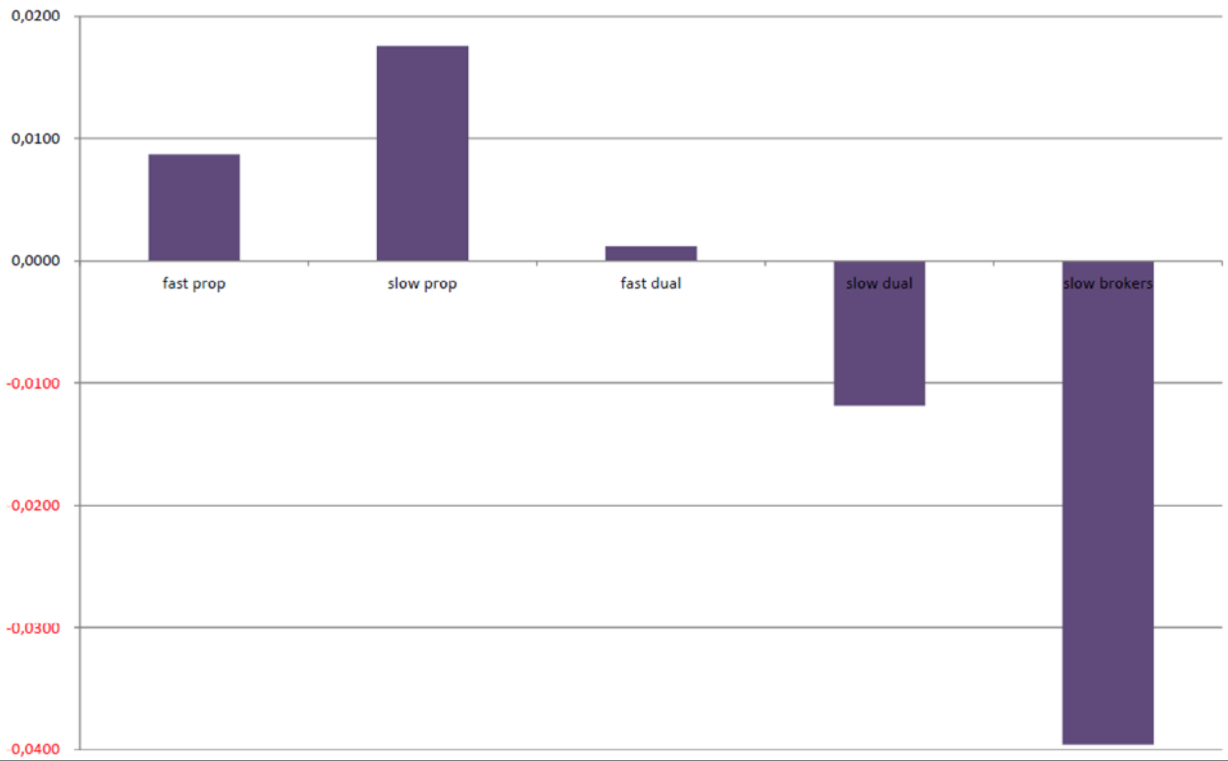

Panel B: Marketable orders' profits by periods and by capitalisation

Graph 8

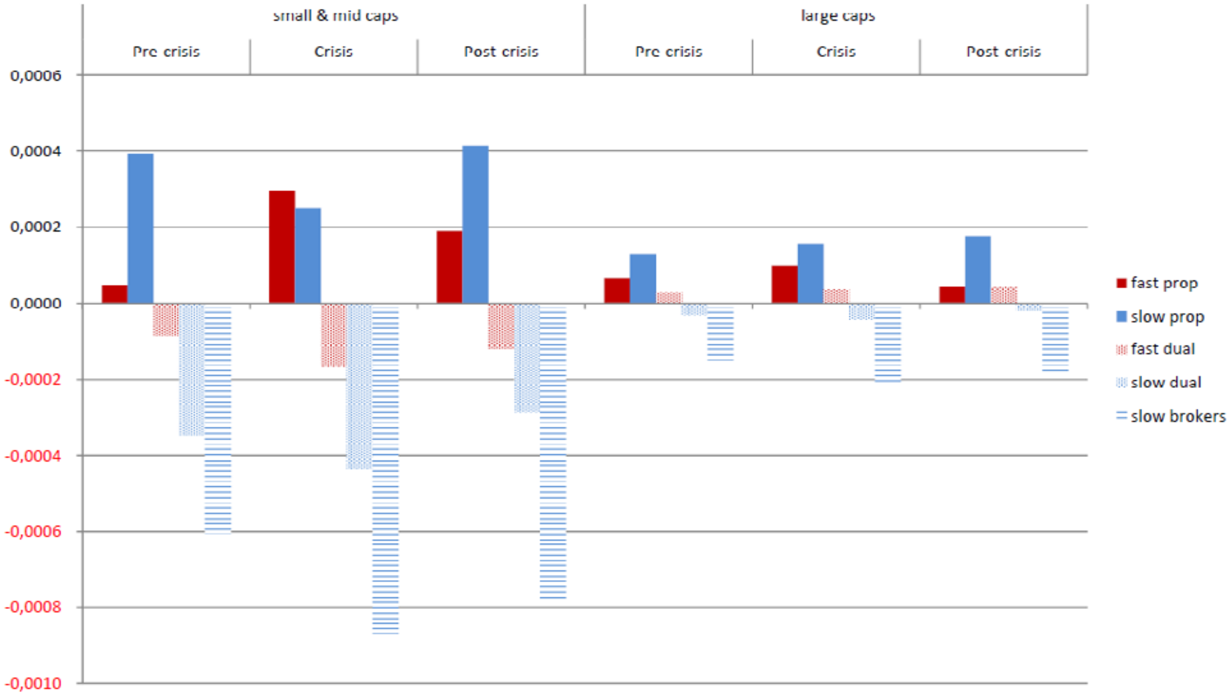

Prob(marketable order *n*|marketable order *n-1*)

Graph 9

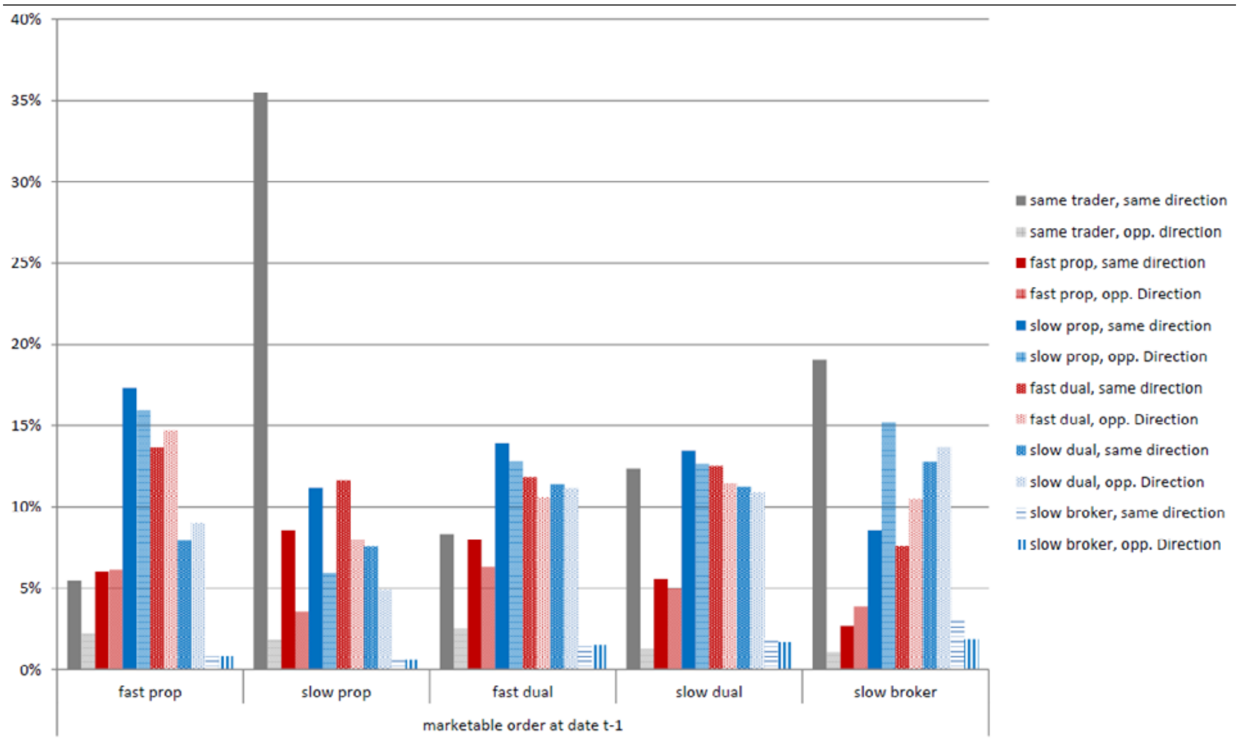

Panel A: Adverse selection cost for (non-immediately executed) limit orders

$(M_{t+2min} – M_{t-})/(M_{t-})$ * (sign of limit order) *100 Graph 10

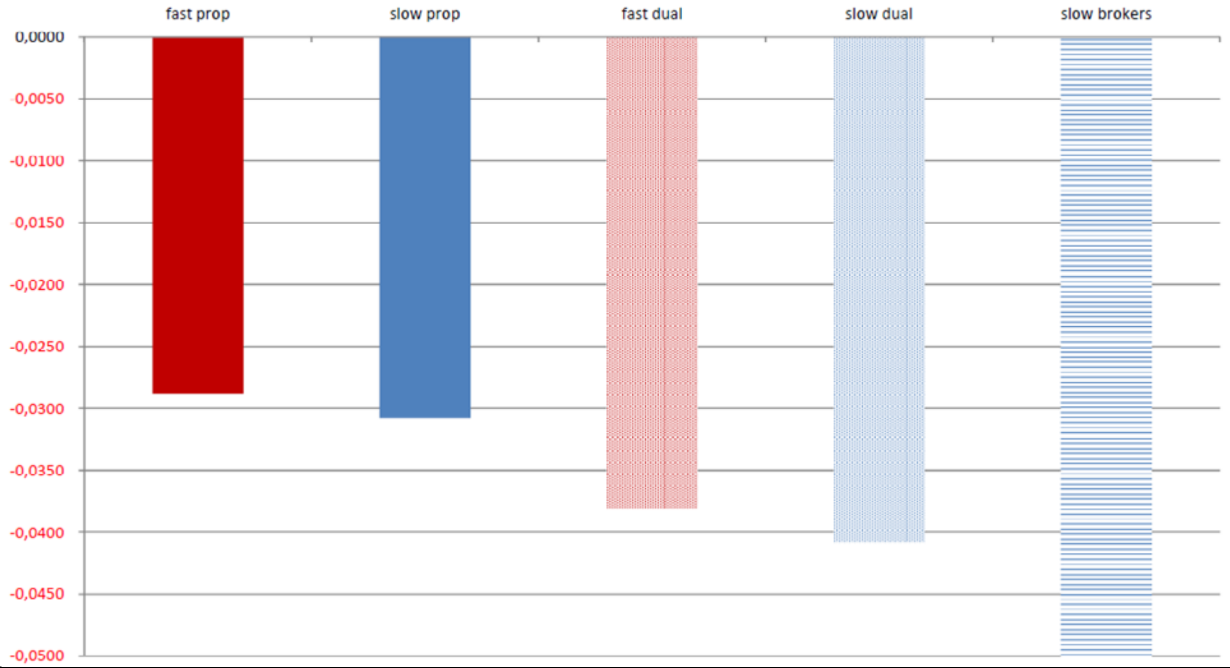

Panel B: Adverse selection cost during crisis Graph 10

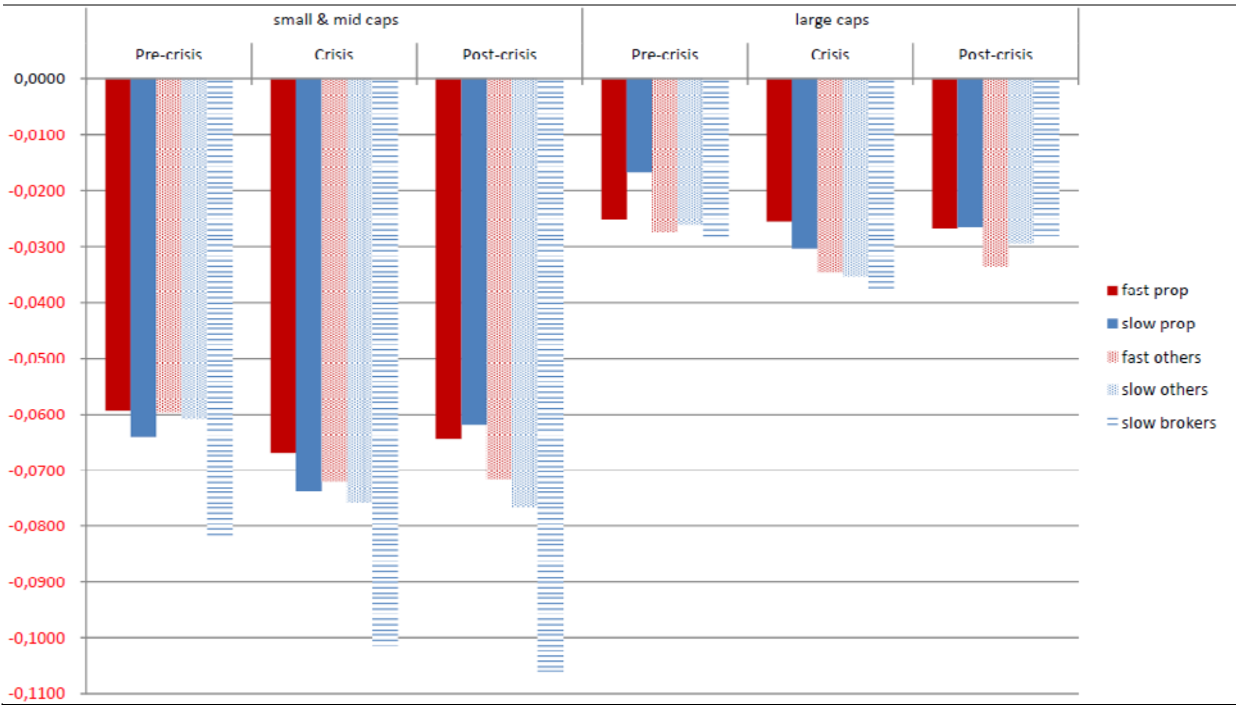

Panel A: Profits of limit orders

(Mt+2min – Pt)/(Mt-) * (sign of limit order) * 100 Graph 11

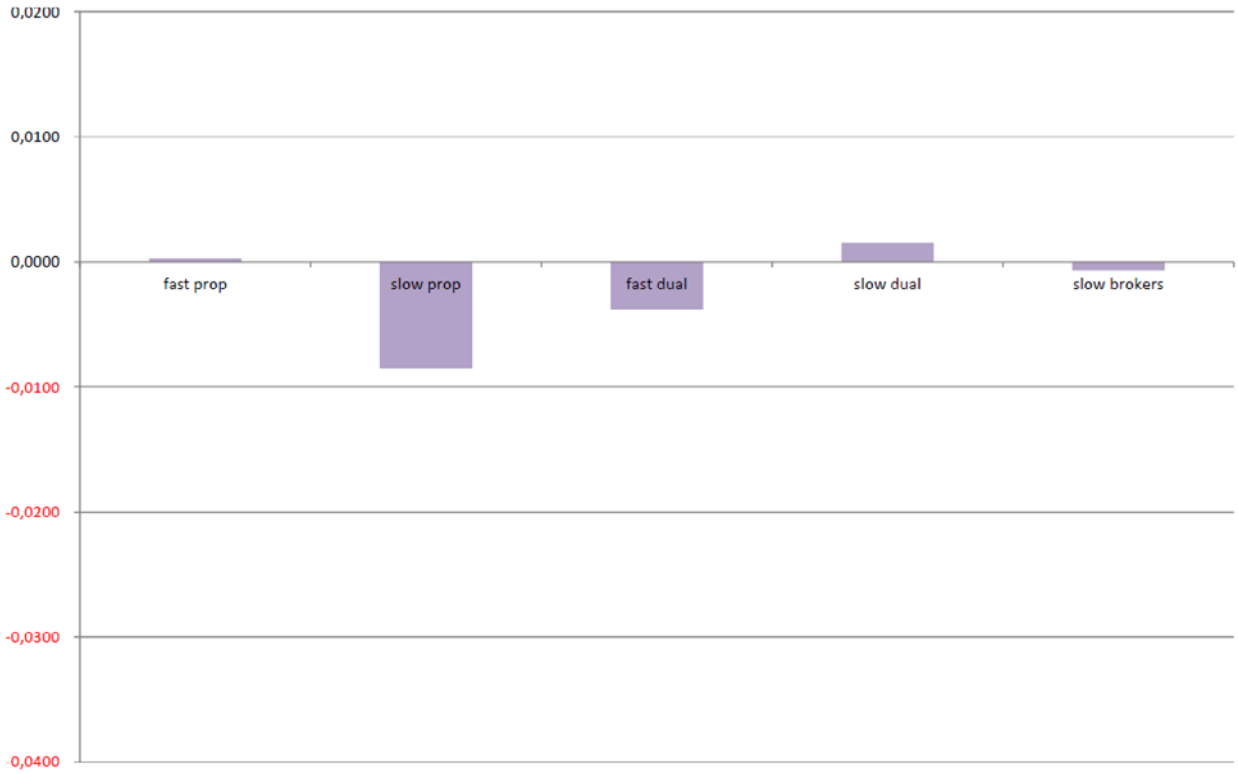

Panel B: Profits of limit orders by capitalisation and period Graph 11

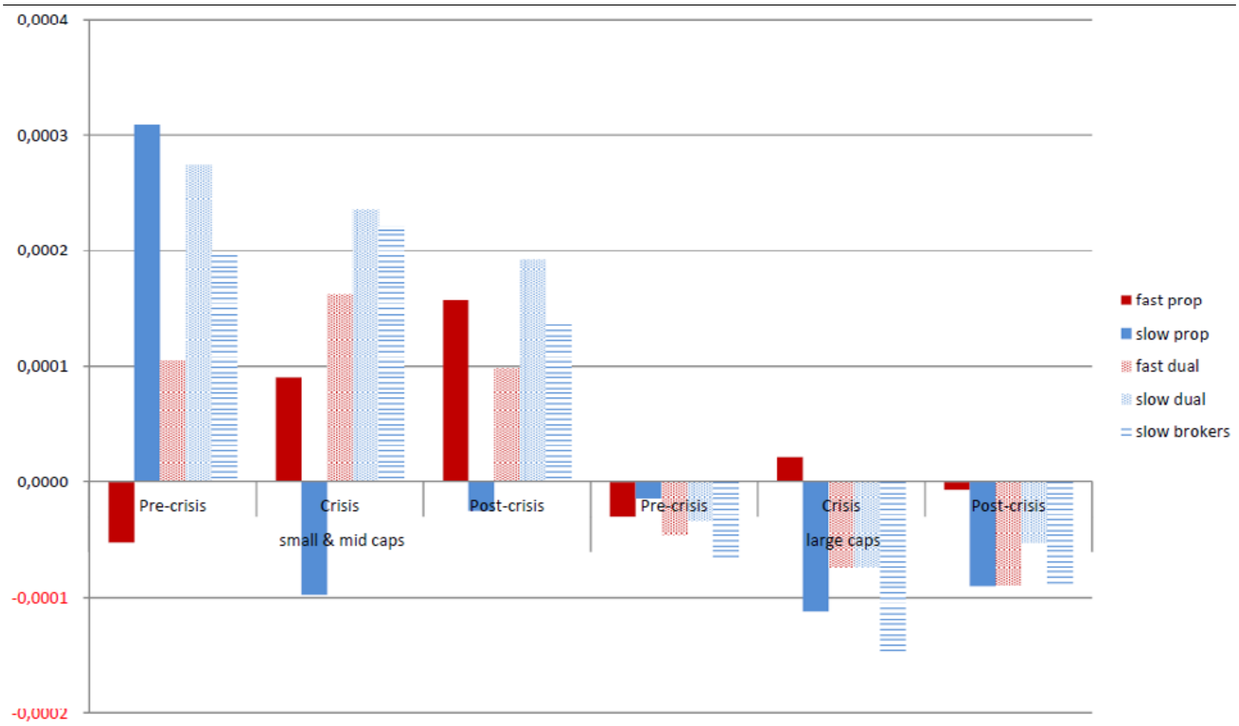

# of cancellations & # of updates to less aggressive/# of trades Graph 12

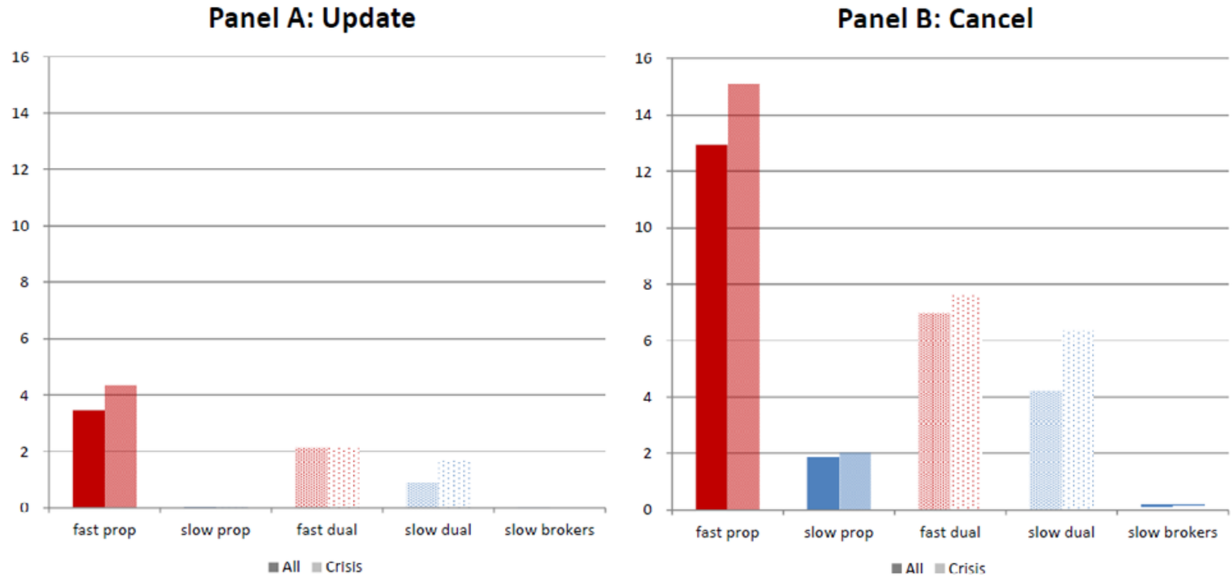

Comments by Arminio Fraga[1]

Liquidity: where are we and two paradoxes

1. Comments on Biais, Declerck and Moinas

Biais et al have written an excellent paper that makes creative use of micro data on exchange-trade stocks. The conclusions are interesting and I have no disagreements to report, only a couple of comments.

First, I tend not to be overly confident in the labelling of various types of investors, and suspect some of the conclusions on the profitability of certain groups may depend on the time horizon that one looks at.

Second, issues related to the speed at which individual participants are able to trade are meaningful, but not entirely new (think of pigeons, the telegraph, the telephone etc). What seems to be more important are asymmetries of information; more on this ahead.

2. Current liquidity in markets is low

Liquidity in equity markets as measured by turnover (ie volume traded divided by market capitalisation) is lower than pre-crisis levels, especially when China is not included in the calculation.

Similarly, liquidity in the markets for government and corporate bonds is also low, some 70% (governments) to 30% (high-yield corporates) down from the highs of recent years. These are all over-the-counter markets.

3. Paradox 1: macro liquidity is high, but micro market liquidity is low

While trading liquidity is lower than recent highs in almost all asset classes and markets, macroeconomic liquidity has never been higher, as indicated by extremely low or negative interest rates in Europe, Japan and the US, as well as by the sizeable quantitative easing (QE) policies in these economies. Moreover, exchange rate intervention in most Asian markets adds up to roughly the same very large amount of asset-buying by central banks as the QE policies of the advanced economies. Why then is market liquidity so low?

For starters, the massive outright accumulation of positions by central banks has driven down availability and probably liquidity as well in government bond markets,

[1] Gávea Investimentos.

especially at the short end of the curve, a crucial source of safe and liquid assets in most economies.

Moreover, higher capital and liquidity rules, as well as restrictions on proprietary trading at banks, have probably contributed to the decline in trading liquidity.

An interesting issue to monitor is whether investors such as large mutual and hedge funds will take over the role of liquidity providers (probably yes) and how far they could go without attracting regulation. For larger issues, such as government bonds, this may well happen on organised exchanges.

The welfare implications of all this require further analytical work, an interesting and challenging task. Some aspects of this question are far from obvious. For example, in pre-crisis days market-makers may have taken on too much risk and benefited from observing the flows of customers. Therefore, it is not clear whether such lower liquidity (driven by less extreme risk-taking and less asymmetric information) is welfare-reducing.

4. Paradox 2: liquidity is only there in quiet times, when not needed – why?

Here the problem may be one of herding, driven by the extreme concentration of volume in a few large central banks and asset managers, which leads to herd behaviour, perhaps in turn driven by a collective focus on some key risk factors such as changes in monetary policy. Markets seem to be more and more characterised by smooth trends that last for a while but are interrupted by violent spikes.

Here the macro may be connected to the micro: the mispricing of liquidity and extremely large aggregate maturity mismatches (such as those that prevailed before the crisis of 2007–08, and may still prevail) may be the product of regulations such as formal and implicit deposit insurance, the availability of lender of last resort support and asymmetric monetary policies. This is a hot research topic, and we have to thank the BIS for all the good work it is doing in the area.

Comments by Francesco Papadia[1]

Let me first of all thank the organisers for the invitation to this interesting conference and for the chance to be back in the stimulating BIS environment.

Then, before discussing this interesting paper by Bruno and colleagues, let me make two initial remarks.

The first is that I accept all the interpretations that Bruno and colleagues give of various price developments around order time (information content, adverse selection, profitability).

The second initial remark is about the definition of liquidity. There are at least three separate, but connected, concepts of liquidity: funding liquidity, central bank liquidity and market liquidity. Bruno's paper is on the latter. In particular, he looks at the liquidity of stocks on the basis of a huge number of observations for 2010. Nowadays, as Jaime Caruana said in his introductory speech, we are more worried about bond rather than stock liquidity, but looking at the latter can shed some light also on the former.

I started reading the interesting and intriguing analysis by Bruno and colleagues looking for answers to six policy questions:

1. Is the complaint of bankers that recent regulatory changes limiting proprietary trading have negatively affected liquidity, because intermediaries are no longer able to buffer liquidity shocks by warehousing securities, justified?
2. Does proprietary trading have a socially valuable function, in particular enhancing market liquidity?
3. Does fast trading have a socially valuable function?
4. Have there been innovations in market organisation – essentially IT developments – over the last few years leading to easier matching of "natural" sellers and buyers, thus reducing the need of securities warehousing?
5. Have there been in recent years innovations in the distribution and processing of information such that adverse selection has been reduced as less weight is given to orders because of better fundamental information?
6. Is more abundant **funding** liquidity, essentially because of the provision of central bank liquidity, helping **market** liquidity?

The paper by Bruno is very useful in shedding light on the first three of these questions.

Let me come to my understanding of what Bruno's results have to say about them.

I think questions 1 and 2 can be answered together, following Bruno's analysis, in the affirmative: regulations constraining proprietary trading negatively affect liquidity and proprietary trading has a positive social value. I draw this conclusion from the following results:

[1] Bruegel.

Proprietary traders, belonging to both the fast and the slow category, place contrarian "take" orders, even during crises, instead of following momentum. Their ability to do so depends on not having (or having fewer) agency problems, as they act as principals, having better risk-bearing capacity and suffering less adverse selection, especially when belonging to the "fast" category. And they are profitable. The latter point is more indirect evidence of a socially useful function, if you remember, by analogy, Friedman's conclusion that only profitable foreign exchange interventions have stabilising properties. Furthermore, from a Darwinian perspective, profitability assures the ability of proprietary traders to perform their stabilising function in a persistent way – obviously if not hindered by regulation.

Let me now come to the third question. Overall in the results provided by Bruno and colleagues I find weaker evidence of a positive social function for fast traders as such. Fast proprietary traders share the general positive characteristics of proprietary traders of being contrarian in their take orders, but this is a characteristic tied to them being proprietary traders rather than fast traders. Of course, they put forward many more orders than slow traders and cancel and modify them much more frequently, but they tend to be, unless they are proprietary traders, momentum traders, thus contributing to sustained price changes. It is true that they appear, in their make orders, less exposed to adverse selection than slow traders, but belonging to the category of proprietary traders is more important in this respect than being fast or slow traders, as one can see in the following table, reproducing some results provided in the paper.

Measure of adverse selection (in basis points) for different categories of traders in marketable trades

Table 1

	Fast	Slow
Proprietary	2.6	3.5
Dual	3.7	4.1

Fast traders are less exposed to adverse selection than slow traders in both proprietary and dual trader categories. But slow proprietary traders are less exposed to it than fast dual traders, showing that belonging to one category is more important than belonging to the other. In addition, overall, fast traders lose money, thus shedding doubts about their ability to continue trading permanently.

These are, if my interpretation is correct, important conclusions one can draw from Bruno's analysis. In their light, measures limiting proprietary trading carry negative implications for liquidity. Instead, the other policy implication drawn by Bruno, ie that capping the ratio of the number of messages to the number of trades by participant could be counterproductive, is not so obvious to me. In any case, the three policy questions that are not addressed by Bruno's analysis need to be answered before drawing final policy implications. Let me be more specific.

If one could conclude that over the last few years there have been innovations in market organisation – essentially IT developments – leading to easier matching of "natural" sellers and buyers, then the fact that securities warehousing is more difficult, because of regulation constraining banks' proprietary trading, would be less of a problem and could be more easily offset by the financial stability advantages of such a regulation. An analogy with "just in time" manufacturing, where there is much less need for warehousing of intermediate goods, is obvious here. The same kind of reasoning can be made on a prospective basis: if the market needs time to adapt its organisation to the weaker presence of warehousing agents, then the disadvantage of lesser proprietary trading is only temporary.

A similar argument can be made about the better distribution and process of information: if prices are more firmly anchored by available fundamental information, the fact that the action of proprietary traders and possibly also that of fast traders, which are less affected by adverse selection, finds an obstacle in regulation is less of a problem. And again, the price of less buffering could be paid more easily for the sake of more financial stability.

The third point I mentioned above is whether better funding liquidity provided by more abundant central bank liquidity can offset the negative effect of the weakening of the stabilising function of proprietary trading. This factor can have some relevance in the short and medium run, but is the least important on a medium- to long-term basis. Indeed one hopes that central banks will not need to maintain their exceptionally large provision of liquidity on a permanent basis.

In conclusion, the three policy issues for which Bruno's paper contains relevant information are important and need discussion. Bruno's paper advances our knowledge of who provides liquidity and raises a question about the effect on liquidity of the new regulatory constraints on banks to carry out proprietary trading. In addition, Bruno's analysis raises questions, in my view, about a clear and strong social value of fast trading. However, before translating these analytical points into policy conclusions, more evidence is needed on other important aspects of market organisation which are not addressed in the analysis.

Previous volumes in this series

No	Title	Author
562 May 2016	Expectations and investment	Nicola Gennaioli, Yueran Ma and Andrei Shleifer
561 May 2016	Mobile collateral versus immobile collateral	Gary Gorton and Tyler Muir
560 May 2016	Has the pricing of stocks become more global?	Ivan Petzev, Andreas Schrimpf and Alexander F Wagner
559 April 2016	A comparative analysis of developments in central bank balance sheet composition	Christiaan Pattipeilohy
558 April 2016	Why bank capital matters for monetary policy	Leonardo Gambacorta and Hyun Song Shin
557 April 2016	How Does Bank Capital Affect the Supply of Mortgages? Evidence from a Randomized Experiment	Valentina Michelangeli and Enrico Sette
556 April 2016	Threat of Entry and Debt Maturity: Evidence from Airlines	Gianpaolo Parise
555 March 2016	The causal effect of house prices on mortgage demand and mortgage supply: evidence from Switzerland	Christoph Basten and Catherine Koch
554 March 2016	Can a Bank Run Be Stopped? Government Guarantees and the Run on Continental Illinois	Mark Carlson and Jonathan Rose
553 March 2016	What drives the short-run costs of fiscal consolidation? Evidence from OECD countries	Ryan Banerjee and Fabrizio Zampolli
552 March 2016	Fiscal sustainability and the financial cycle	Claudio Borio, Marco Lombardi and Fabrizio Zampolli
551 March 2016	When the Walk is not Random: Commodity Prices and Exchange Rates	Emanuel Kohlscheen, Fernando H. Avalos and Andreas Schrimpf
550 March 2016	A new dimension to currency mismatches in the emerging markets: non-financial companies	Michael Chui, Emese Kuruc and Philip Turner
549 March 2016	Monetary policy spillovers and currency networks in cross-border bank lending	Stefan Avdjiev and Előd Takáts
548 March 2016	Moving in tandem: bank provisioning in emerging market economies	Andres Murcia and Emanuel Kohlscheen
547 March 2016	When pegging ties your hands Intertemporal considerations in currency crises	Nikola Tarashev and Anna Zabai
546 March 2016	Financial intermediation and monetary policy transmission in EMEs: What has changed post-2008 crisis?	M.S. Mohanty and Kumar Rishabh

All volumes are available on our website www.bis.org.