Anonymity in a Limit-Order Market: An Experimental Analysis

Christophe Majois*

January 2008

Abstract

In this paper, we adopt an experimental approach to evaluate the impact of pre-trade anonymity in order-driven markets. More specifically, we build an experimental design of an electronic limit-order market, and compare two settings: one in which traders observe the identities of agents placing orders in the order book, and the other one where this information is not available on market screens. We find that ID code disclosure does not alter liquidity, efficiency or traders' profits. This is in contrast to the finding of Perotti & Rindi (2006) that anonymity improves market liquidity in a design where information acquisition is endogenous. When we exogenously increase the proportion of informed traders in our experimental markets, anonymity remains non-significant but we observe a clear improvement in liquidity and efficiency. Our experiments thus cast doubt on the real role played by anonymity in stock markets.

JEL Classification: G14, G10
EFM Classification: 360
Keywords: Market Microstructure, Anonymity, Experiments

*FNRS Research Fellow, Louvain School of Management and FUCaM. I wish to thank Christophe Bisière, Robert Bloomfield, Rudy De Winne, Hans Degryse, Carole Gresse, Maureen O'Hara, Isabelle Platten and Gideon Saar for helpful comments and suggestions. I also thank the ICM Doctoral Fellowship program for financial support. Contact address: FUCaM (Catholic University of Mons), Chaussée de Binche 151, 7000 Mons, Belgium. Phone: +32 65 323 336. Fax: +32 65 323 223. E-Mail: christophe.majois@fucam.ac.be.

1 Introduction

What precise information should stock exchanges display to market participants? The answer to that question helps define the level of transparency that markets offer to their users. Transparency in microstructure is defined by O'Hara (1995) as the "*ability of market participants to observe the information in the trading process*", i.e., information related to order flow, identities of market participants, price and volume of past trades, etc. Transparency is usually divided into two types: (i) *pre-trade transparency* refers to the ability to observe the order flow (in an order-driven market) or the quotes of market makers (in a quote-driven market), and the identities of traders who post these orders or quotes; (ii) *post-trade transparency* refers to the revelation of information relative to past transactions: quantity, prices and traders' identities.

Many academic papers have addressed the question of whether there exists an optimal transparency level in financial markets. While the common intuition is that transparency should benefit market quality, the results from the literature sketch a far more complex story.[1] Stock exchange themselves often modify their trading rules regarding their transparency level.

In this paper, we are interested in a very specific element of pre-trade transparency, namely the issue of anonymity in an order-driven market. The question we address is the following: what are the advantages and drawbacks of displaying the ID codes of liquidity providers in an electronic order book? This question is interesting as several stock exchanges have modified the level of anonymity in recent years. While the Korean Stock Exchange switched from an anonymous to a non-anonymous market in 1999, Euronext Paris, the Tokyo Stock Exchange and the Australian Stock Exchange have experienced the opposite move in 2001, 2003 and 2005 respectively. The Toronto Stock Exchange also launched the "attribution choice" in late 2005, i.e., a feature that allows traders to display or not their ID code alongside their orders. So the latest trend is clearly toward more anonymous systems, but it remains to be seen if this is indeed an optimal move from a market quality perspective.

Working with field data to deal with this issue is often problematic, given that there are numerous sources of noise in the real world that can pollute the results and make it difficult

[1] See for instance the debate on order book disclosure. On one hand, Madhavan et al. (2005) show that the real-time public dissemination of depth and quote at the five best limits on the Toronto Stock Exchange in 1990 has been followed by a significant increase in quoted and effective spreads. Boehmer et al. (2005) on the other hand find that the introduction of OpenBook on the New York Stock Exchange in January 2002, which has enabled traders off the exchange floor to observe depth in the limit order book in real time, has improved price efficiency and liquidity, so they conclude that transparency is a "win-win" situation that benefits investors.

to derive any definitive conclusion.[2] That is why we have decided to adopt an experimental approach to study this issue. Experimentation is an interesting research method, as it provides the researchers with a level of control that is unattainable when working with field data, which also helps to establish a clear causality link between the treatment variable and the observed differences. Another advantage of the experimental method is that some variables that are observed only with much difficulty in the real world can be more easily studied in a laboratory market. It is for instance rather hard to evaluate how efficient real markets are, given the difficulty to assess the true value of assets. This problem vanishes in an experimental setting, as the true value of the security is known for sure to the experimenter.

Other questions relative to transparency in financial markets have been addressed through laboratory experiments. For instance, Bloomfield & O'Hara (1999) analyze the influence of quote and trade transparency in a noncontinuous quote-driven market, and show that trade disclosure improves market efficiency, hurts liquidity (spreads are wider at the beginning of the game in a transparent setting) and benefits market makers at the expense of informed traders, while quote transparency does not influence market performance. By contrast, Flood et al. (1999) show that quote transparency in a continuous multiple-dealer market generates a clear trade-off between liquidity and market efficiency: transparent markets are more liquid but less informationally efficient in their experimental setting.

Experimentation has already been adopted by Perotti & Rindi (2006) in order to study the anonymity issue. These authors implement a laboratory experiment of a continuous order-driven market, where traders can chooser whether they want to be informed or not. Within this particular design, they show that anonymity induces more traders to acquire information, and so improves liquidity given that informed traders are the best liquidity providers. Anonymity does not have any influence on market efficiency.

Our experiment uses a more classical design, where the number of informed traders is fixed exogenously and is kept constant all the time. The only difference between structures lies in the information that is displayed on traders' screen. Changes in the dependent variables are thus clearly attributed only to the differing information concerning the identity of traders in the order book.

Our results show that anonymity does not have any impact on bid-ask spreads, absolute pricing errors and traders' profits. In digging for why this "no-effect" result holds, we look at the behavior of traders in our experimental markets. It turns out that informed traders act

[2] Majois (2007), for instance, focuses on the switch to anonymity on Euronext Paris in April 2001 and argues that results previously found in the literature may be attributable to a global liquidity phenomenon

aggressively at the beginning of the trading game, and compete with one another in order to reap the profit opportunities provided by uninformed traders. Toward the end of the game, i.e., when the market has discovered the true security value, informed traders start to behave as liquidity providers, thus revealing themselves in the order book. However, this identity disclosure does not bring any "new" information as the game is somehow over.

The result on bid-ask spreads contradicts the improvement in liquidity found by Perotti & Rindi (2006). More precisely, the authors show that the increase in the number of informed traders in anonymous markets is directly related to the decrease in spreads. That is why we have implemented a second design where we have *exogenously* increased the proportion of informed traders. Both liquidity and efficiency are significantly improved in the second set of experiments relative to the first design – but anonymity still does not play any role. Those results cast doubt on the exact role played by anonymity in Perotti & Rindi's (2006) experiments.

This paper is organized as follows. In Section 2, we provide a review of the market microstructure literature related to anonymity. Section 3 develops our research hypotheses. The experimental design is detailed in Section 4 and results are provided in the fifth section. Section 6 contains the results for the second set of experiments. The last section concludes.

2 Anonymity in the Literature

As any other issue in market transparency, anonymity can be considered both pre-trade and post-trade.

The issue of post-trade anonymity has attracted a relatively low volume of research. Fishman & Hagerty (1995) develop a two-period model where the transactions of an insider at the first period can be displayed to market participants before the second trading period. The authors demonstrate that this setting may benefit the insider at the expense of uninformed traders, as the former can manipulate the market in order to get a profit during the second trading period. Frino et al. (2005) develop a theoretical model, similar to Glosten & Milgrom (1985) and Easley & O'Hara (1987), and show that two consecutive transactions in the same direction that originate from the same trader have a higher informational content than if they are initiated by two different traders. Consistent with this hypothesis, they find that the price impact on the Australian Stock Exchange (ASX) is bigger for transactions initiated by the same trader than for transactions initiated by two different brokers. As the ASX displays post-trade ID codes to market participants, the authors conclude that post-trade identity disclosure benefits market efficiency. The positive view of post-trade ID disclosure is shared by two other papers. Eom et al. (2006) analyze data from the Korea Exchange, where the largest five cumulative sellers and buyers of the day are continuously disclosed. Among other results, it appears that liquidity increases and returns become more negative (positive) when foreign securities firms are added to the list of large sellers (buyers). This indicates not only that foreign firms have superior information, but also that this information is disseminated to the market. Waisburd (2004) examines the effect of the disclosure of post-trade identities on market liquidity, using 27 stocks from the Paris Bourse that changed market category, between the *anonymous* "Continu A" and the *non-anonymous* "Continu B". He concludes that post-trade identity disclosure reduces execution costs, and hypothesizes that this improvement in liquidity mainly results from a reduction in inventory control costs, as adverse selection does not seem to be affected by post-trade identity revelation.

By contrast, pre-trade anonymity has been more extensively analyzed. It may be convenient to subdivide this issue in three topics. The first one deals with the identity of *liquidity demanders*. Admati & Pfleiderer (1991) theoretically study the effects of the practice of sunshine trading, i.e., the preannouncement of trades used by some liquidity traders. This preannouncement identifies the trader as informationless, thus reducing its adverse selection costs. That is why the practice of sunshine trading has positive effects on liquidity traders who engage in such preannouncements. Dia & Pouget (2006) provide an empirical support in favor of the existence of sunshine trading. Using data from the West-African Bourse, they show that large orders that are placed early during the preopening period and are not canceled, are traded without significant price movements. Madhavan (1996) develops a model where the order flow is divided into (i) a price-sensitive component which arises from the strategic decisions of speculative traders and liquidity providers, and (ii) a price-inelastic component that arises from the liquidity demand of traders and is by definition not related to information about the asset fundamental value. Disclosing the price-inelastic component of the order flow to market participants prior to trading (e.g., thanks to sunshine trading) leads to more informative prices, but the impact on price volatility and market liquidity is less clear. Fishman & Longstaff (1992) and Roëll (1990) consider the dual trading practice, where brokers can trade both for their own account and for their clients. As brokers usually know their clients, they can infer their trading motivations. In both models, dual trading clearly benefits uninformed clients.

Forster & George (1992) argue that the anonymous execution of liquidity-motivated trades implies that neither the direction nor the magnitude of liquidity trades is known. They analyze the effects of revealing this direction and this magnitude separately. Summarized briefly, their results tend to show that less anonymity has no impact on market depth or price efficiency, while it seems to profit liquidity traders provided that the revelation is made to all market participants. Benveniste et al. (1992) develop a theoretical model where brokers act on behalf of their clients, who themselves trade either for liquidity needs or because they have some private information. All trades are made with a risk-neutral specialist. The authors show that in a market setting where the specialist may observe – after the trade – whatever information was available to the broker at the time the trade was executed a separating equilibrium is attainable, where the specialist sets different spreads depending on the motivation of the trades, and which Pareto-dominates a pooling equilibrium. Benveniste et al. (1992) conclude that the specialist system can be viewed as a market mechanism that improves the welfare of exchange members and the terms of trade for public customers by reducing the incentives to exploit private information. Analyzing situations where the specialist physically move on the NYSE floor, Battalio et al. (2007) provide a direct empirical support for the theoretical predictions in Benveniste et al. (1992). In the same spirit, Garfinkel & Nimalendran find that NYSE market makers are better able to identify the probably informed trades, so they conclude that this evidence is consistent with less anonymity in the NYSE specialist system compared to the NASDAQ dealer system. Seppi's (1990) analysis focuses on equilibrium in block trading. His conclusion is that the non-anonymous upstairs market, i.e., block trading that does not occur on the main exchange, is primarily used by uninformed liquidity traders. This result is driven by the possibility offered to the liquidity trader to be identified as being uninformed when trading in the upstairs market. Madhavan & Cheng (1997), Smith et al. (2001), Booth et al. (2002) and Bessembinder & Venkataraman (2004) provide empirical evidence in support of this prediction on different markets.

In summary, most papers dealing with the anonymity of liquidity demanders, with notable exception of Madhavan (1996), agree upon the idea that revealing the identity of uninformed liquidity traders should improve the terms of the trades for them.

A second strand of literature focuses on the case of *parallel markets*, when a single security can be traded on different trading structures, and those structures vary in the level of anonymity. Desgranges & Foucault (2005) develop a model in the spirit of Benveniste et al. (1992), where an investor can trade either with a dealer with whom he develops a long-term relationship, or by placing an order on an anonymous market where dealers are competing with each other. The authors show that the practice of price improvement emerges as a result of this long-term relationships. Dealers cream-skin uninformed transactions, and this increases the level of information asymmetry on the anonymous market, where spreads become higher. The long-term relationships thus negatively affect those investors who cannot engage in such relationships and only trade in the anonymous structure. Theissen's (2003) results empirically support the price improvement hypothesis. Using data from the Frankfurt Stock Exchange, he finds that non-anonymity allows the specialist to assess the probability that a trader is informed and that he uses this knowledge to price-discriminate, by quoting a large spread and granting price improvement to traders deemed uninformed.

The competition between Nasdaq market makers and the anonymous Electronic Communication Networks (ECNs) is at the center of several articles. Barclay et al. (2003) show that ECNs attract more informed orders compared to Nasdaq market makers – the permanent price impact is higher for ECN trades than for market maker trades, and ECN trading explains 60 to 100 % more of the stock price variance than market maker trading. They interpret this result as a direct consequence of the anonymous trading provided by ECNs. Simaan et al. (2003) analyze the quoting behavior of market makers on Nasdaq and ECNs, and show that they quote more aggressively (narrow the spread) when they post limit orders through the ECNs. Their interpretation is that an anonymous system deters the emergence of a tacit collusion outcome, and so a lower level of pre-trade transparency should help improve competition on the Nasdaq. The German market also presents features of parallel structures. Grammig et al. (2001) compute the probability of information-based trading (PIN) for 30 stocks of the DAX index traded on the non-anonymous floor of the Frankfurt Stock Exchange and on the anonymous screen-trading system. Given that the PIN is significantly lower on the floor than in the screen-trading environment, they conclude that informed traders indeed prefer to trade in an anonymous electronic system. Similarly, Theissen (2002) compares transaction costs for stocks simultaneously traded on the floor of the Frankfurt Stock Exchange and the IBIS electronic screen-trading system, and finds that the adverse selection component of the spread is lower on the floor.[3] Jain et al. (2006) focus their analysis on the London Stock Exchange, where an anonymous order book (SETS) operates alongside a non-anonymous dealer market. While the PIN does not differ across the systems, the permanent price impact is higher on SETS than on the dealer market. All the above papers agree that an anonymous system attracts more informed traders. To our knowledge, the only counter-evidence to this widely accepted fact is provided in Reiss & Werner (2005), who also analyze the London Stock Exchange but

[3] Heidle & Huang (2002) also find that the PIN and spreads decrease when firms relocate from a "dealer" market (Nasdaq) to an auction market (NYSE or Amex), and interpret this finding as related to the level of anonymity in both systems.

focus on inter-dealer trades. Surprisingly, they show that adverse selection is more present in the direct non-anonymous public market than in the anonymous brokered systems. The focus on inter-dealer trades and the fact that only dealers could access the brokered systems most probably account for the differing results.

The third topic deals with the identity of *liquidity providers*. Foucault et al. (2007) develop a limit order model where some liquidity providers have information about future volatility – this information is valuable given the option-like feature of limit orders. In this model, the spread provides a signal about future volatility. The authors demonstrate that anonymity should influence the spread and its informational content in the same way, but the precise direction (increase or decrease) depends on a specific model parameter. They then use the natural experiment provided by the switch to anonymity on Euronext Paris on April 23, 2001, and show that both the spread and its informational content have decreased after that date. Rindi (2004) studies the effects of pre-trade anonymity on market quality in an order-driven setting. She considers a batch auction where all orders posted by risk-averse informed, uninformed, and noise (liquidity) traders, are cleared at a single price, in the spirit of Grossman & Stiglitz (1980). The analysis focuses on two extreme transparency regimes: an anonymous market where traders only observe the current price, and a fully transparent market where traders observe the order flow, the clearing price and the personal identities.[4] When there is a given proportion of informed agents, transparency enhances liquidity, which is measured by the price impact of a noise trader's order. However, the opposite result holds when the number of informed traders is determined endogenously. The key element driving those results is that informed traders in this model are the best liquidity providers. When the number of informed traders is fixed exogenously, transparency makes the uninformed quasi-informed, thus enhancing quality. When information acquisition is endogenous, transparency reduces the incentive to acquire information, leading to a lower number of informed traders and to less liquidity in the market. In this model, informational efficiency is always made better by an increase in transparency, while the effect on volatility is unclear.

On the empirical side, Comerton-Forde et al. (2005) use natural experiments that took place on Euronext Paris, the Korean Stock Exchange and the Tokyo Stock Exchange and provide further evidence that anonymity improves liquidity in order-driven markets. Two papers have analyzed the switch to anonymity on the Italian secondary bond market (MTS) in 1997. Focusing on investors' behavior, Albanesi & Rindi (2000) show that there has been an

[4]A third, partially transparent regime, where traders observe the order flow and the clearing price but not the identities, is also taken into account but less fully analyzed.

important decrease in the positive autocorrelation of trades after the introduction of anonymity. They attribute this finding to a decrease in order fragmentation by market makers, as there is less risk of free-riding by small traders when the identity is not displayed. Scalia & Vacca (1999) show that anonymity has made liquidity traders worse off, and large/informed traders better off. It has also improved liquidity through an increase in the number of active bonds and a decrease in bid-ask spreads, while volatility has been reduced.

At the experimental level, the anonymity issue has been addressed indirectly by Plott & Sunder (1982). Their primary objective is to determine which information dissemination model best reflects the results observed in laboratory settings, and the rational expectations model is found to be the winner. This means that agents use market information to update their beliefs on the state of nature that has been randomly drawn. When discussing which precise information is used by traders, Plott & Sunder minimize the possible role played by the identity of traders – their experiments were organized as oral double auctions where subjects could see who were the traders placing orders – and argue that critical bid ans ask probably convey most information. Anonymity is the central focus of Perotti & Rindi (2006), who experimentally test Rindi's (2004) theoretical predictions. The authors construct an automated continuous double auction market and compare the outcomes of a transparent and an anonymous setting. In order to endogenize the acquisition of information, each game is preceded by a "market for information" where traders can buy a signal about the asset liquidation value. While the number of informed traders is known in the anonymous setting, the number *and* the identity of informed traders are known to everybody in the transparent market. Consistent with theory, transparency reduces the incentive to acquire information and decreases liquidity and volatility. The results on informational efficiency are less clear, and do not confirm the theoretical hypothesis that transparency would improve efficiency.

Our paper is closely related to Perotti & Rindi's (2006) analysis, and it is worthwhile to highlight in which respect our experiment differs from their design. A key element in Perotti & Rindi's (2006) article is that they test the version of Rindi's (2004) theoretical model where information acquisition is endogenous, i.e., they introduce in their experimental design a market for information during which traders can choose whether or not they buy information on the asset value. However, as described above, Rindi (2004) also develops a more traditional model version where the number of informed traders is constant, and that case is worth analyzing for at least three reasons.

Firstly, the theoretical predictions on the impact of anonymity on liquidity crucially depend on the version under consideration. When information acquisition is endogenous, transparency reduces the incentive to acquire information and, given that informed traders are the best liquidity providers, a reduction in their number hurts liquidity. This is exactly the result obtained experimentally in Perotti & Rindi (2006). However, the opposite effect appears when the number of informed traders is constant: in that case, transparency makes the uninformed quasi-informed, which improves market liquidity. It is interesting to see whether that alternative prediction also hold in a laboratory setting.

Secondly, the endogenous information case may not be completely relevant when a single market switches from a transparent to an anonymous system, as it happened for instance on Euronext Paris or the Tokyo Stock Exchange. Rindi (2004) argues that "Big broker-dealers assemble a large number of customers and therefore have access to information on large order flows.[. . .] Large liquidity suppliers, forced by transparency to share their information on order flows, may leave the market." The argument that traders would change their behavior in terms of information gathering because one small piece of information – traders' ID codes – appears (disappears) on (from) the market screen is questionable. So focusing on the exogenous case does not only reflect an academic interest, but it is also relevant from a practical point of view.

Finally, the design of the market for information itself may be highly critical. Indeed, Sunder (1992) introduces a market for information before conducting experiments similar to Plott & Sunder (1982), and shows that how the market for information is organized has an impact on the outcomes of the trading markets. More precisely, when the information market is organized with a *fixed quantity* – i.e., there is a limited number of agents who can buy the information, and only the highest bidders get it – the predictions of the full rational expectations equilibrium are observed. But when the market for information is organized with a *fixed price* at which any trader can buy the information, as is the case in Perotti & Rindi's (2006) article, the outcomes of the trading game are better described by the *noisy* rational expectations equilibrium. So an experimental design using a market for information should probably control for the way this market is organized, by adding it as another treatment variable.

3 Hypotheses development

Our experiment deals with the third strand of the literature presented in the above section, that is, we focus on pre-trade anonymity of liquidity providers in an order-driven market. While Foucault et al. (2007) propose a model that more closely applies to the kind of electronic order-driven market we implement in the experiment, we do not use it as our primary source of hypotheses because their results depend on some specific assumptions. For instance, informed liquidity providers in that model only have the information about the probability of occurrence of an event, and not on the direction of the event. The game structure they analyze is also very detailed, and we don't want to impose as much structure in our experiment. Finally, they do not consider the impact that anonymity may have on informational efficiency. So we base our hypotheses on Rindi's (2004) theoretical paper, even if this paper is set in a batch-auction system. This is the same paper that lies at the center of Perotti & Rindi's (2006) experimental analysis, but we consider the model version where there is a constant number of informed traders for reasons that have been explained earlier.

In Rindi's (2004) model, uninformed and informed risk-averse traders provide liquidity to a pool of liquidity traders. Transparency, defined as the disclosure of ID codes, improves market liquidity, which is measured by the inverse of the price impact of a liquidity trader's order. The intuition behind this result is the following. Consider the case where a liquidity trader places a buy order, which causes an initial price increase. If the market is anonymous, an uninformed trader may think that the order originates from an informed trader, hence revises upward her estimate of the value of the asset and places a buy order accordingly, which increases the price further. This misinterpretation does not happen when the market is transparent, as all traders see that the initial order has been placed by a liquidity trader. So the first hypothesis is the following:

Hypothesis 1 A transparent (that is, non-anonymous) market is more liquid than an opaque (that is, anonymous) market.

Informational efficiency, measured as the inverse of the conditional variance of the liquidation value of the asset, is shown to be higher when the market is transparent, which directly follows from the uninformed being quasi-informed thanks to information disclosure. This gives our second hypothesis:

Hypothesis 2 A transparent market is more informationally efficient than an opaque market.

Rindi (2004) also considers the effect of transparency on the welfare of the different types of traders. As liquidity is enhanced by transparency (Hypothesis 1), liquidity traders are better off under a transparent regime. Logically, transparency decreases informed traders' welfare. The third hypothesis sounds as follows:

Hypothesis 3 Liquidity traders (informed traders) are better (worse) off under a transparent market than under an anonymous market.

4 The experimental design

The design implemented in this paper presents many similarities with the experimental design used in Bloomfield et al. (2005). These authors study how different types of traders provide and take liquidity in an asset market organized as an electronic double-auction system. The market in which subjects trade can be seen as a simplified, but still representative version of the trading platform used on many stock exchanges around the world (such as the Toronto Stock Exchange, the Australian Stock Exchange, Euronext, etc.). Our interest lies in the effect of differing transparency levels, so we add a new treatment variable to Bloomfield et al.'s (2005) design.

It is useful to define some conceptual terms that will be used throughout the paper. A *cohort* is a group of subjects who always play together. A *security* is a claim on a terminal dividend, and is identified by its true value. A *market* (or *period*) is a game during which a cohort trades a specific security. A *session* is an experiment during which a given cohort trades in several markets. The treatment variable of the experiment is the *transparency regime*, which is either *anonymous* or *transparent.* The monetary unit employed during the experiment is called the Experiment Currency Unit (ECU).

4.1 Organization of a session

A session lasts for about two hours and fifteen minutes, and involves a cohort of nine subjects. The subjects first read the instructions, which are then reviewed together with the instructor.[5] The participants can ask questions about the trading game, but questions related to specific strategies are not answered. The subjects then participate to four "training" markets (two in each transparency regime), and eighteen "experimental markets" (nine for each transparency regime). The transparency regime is modified every three markets – if the first three markets are transparent, then markets 4 to 6 are anonymous, markets 7 to 9 are transparent, and so forth. In order to average out learning effects, half of the cohorts start with the transparent markets, and the other half with the anonymous setting.

4.2 Organization of a market

In each market, subjects have the opportunity to trade a security, whose value is drawn from a uniform distribution between 1 and 100 ECU, inclusive. The draw is independent from one

[5] An English translation of the instructions distributed to participants is provided in the Appendix.

market to another, so each market can be seen as the trading of a new security.

A market involves three types of traders:

Informed traders are told the true value of the security before trading starts.

Traders with a "buy" objective are asked to complete at least 5 buy transactions during the market; we call them *buyers* in the following.

Traders with a "sell" objective are asked to complete at least 5 sell transactions during the market; we call them *sellers* in the following.

Buyers and sellers can be seen as a proxy for liquidity traders, i.e., traders who face exogenous liquidity needs. Our setting does not include *noise traders*, i.e., agents who have no particular information nor any assigned objective. Indeed, in a recent experimental paper, Bloomfield et al. (2006) show that these traders may harm market efficiency as they drive prices away from their fundamental values, an effect that is stronger when prices are farther away from their value. Including them may thus induce some unnecessary noise in the market, which would make it more difficult to find effects only due to the transparency levels.

There are three traders of each type in each market. Participants are told their role before trading begins. They change role every period in such a way that, at the end of the session, each subject has played six times in each of the three roles (three times in each transparency regime). Each player starts each market with no endowment in cash nor in security. This means that inventory and cash position at the end of a period are not carried forward to the next period. There is no cash nor short-sell constraint.

Traders with an objective incur a penalty of 20 ECU for each transaction that is not completed. The objective pertains to the number of trades completed, and not to the inventory at the end of a market. These liquidity traders can see on their screen the number of transactions they still have to make in order to fulfill their objective.

Each market lasts for 150 seconds, and is organized as an electronic order-driven market, where players may take the following actions:

- place limit orders
- cancel limit orders
- place market orders

The price grid is 1 ECU, and the minimum (maximum) price that can be entered is 1 (100) ECU. Marketable limit orders are not allowed: traders cannot place a buy (sell) order with a limit price equal to or above (below) the current best ask (bid). The only way a trader can make a transaction is by placing a market order, i.e., by accepting the current best bid or ask. All orders are for a unit quantity.

In their experiment, Bloomfield et al. (2005) introduce a pre-trading period, during which traders can place orders. At the end of this pre-trading period, all orders that cross disappear from the order book, in such a way that the order book is non-empty at the beginning of the trading period. We do not include any pre-trading period, so the book is empty when a market begins.

As mentioned above, the transparency regime is the treatment variable, that can take two values. In both transparency regimes, traders can see all orders standing in the limit order book. In the *transparent* regime, traders also see the ID code besides the orders in the order book, and they know which traders are informed. This is the only difference with the *anonymous* regime. In both transparency regimes, post-trade transparency is kept at the same level: traders can see the prices of all past transactions, without any ID code.

The trading screen also displays information that allows each subject to know at any time his cash and share positions, the prices at which he has previously bought and sold and the time remaining until the end of the market.

The experiment was programmed and conducted with the software z-Tree [Fischbacher (2007)]. Figure 1 is an example of the screen that participants face when they trade in a transparent regime.

4.3 Subjects and incentives

Subjects were students at the Catholic University of Mons (Belgium) who took a *Financial Markets* or a *Portfolio Management* class. They were either in third year of a Bachelor program or in first year of a Master program. They had never participated to experiments before.

A subject's gain in ECU for a particular market was computed as the sum of his portfolio value and cash value at the end of the period, less the penalty (if applicable).

The gain in ECU at the end of the experiment was the sum of the gains in every experimental period. This profit in ECU was then converted in euros in such a way that the minimum payment was 10 euros and the maximum payment 30 euros. Subjects were paid by wire transfer.

A possible concern with that incentive scheme is that, when subjects suffer losses during the course of a session, they may be willing to take excessive risk. In order to minimize that kind of gaming behavior, we told subjects that suffering losses does not mean that they will get the minimum payment, as other subjects may have greater losses in absolute value. So their best interest is to play as best as they can during each market.

5 Results

We now turn to the experimental results. Thirteen sessions have been conducted, for a total of 117 participants (9 per session). In the following analyses, we exclude data from the four training markets and only consider the eighteen experimental markets of each session. This means that results are obtained from 234 different securities, 117 in each transparency regime.

5.1 Liquidity

Hypothesis 1 posits that an anonymous market is less liquid than a transparent market. In order to study liquidity, we focus on the bid-ask spread. Each market is divided into 10 intervals of fifteen seconds. Figure 2 shows the evolution of the quoted spread during the trading game, with a distinction between anonymous and transparent markets. The value for each interval is the average across markets of the last quoted spread observed for that interval.

This figure displays a pattern that is consistent with previous results in the literature – see for instance Bloomfield et al. (2005). The spread starts at a relatively large level and then progressively and monotonically decreases throughout the game. No striking difference appears between both transparency regimes. However, this graph may be misleading, as other variables that are not represented could also influence the level of the spread. So we propose to analyze liquidity in a multivariate framework. The regression that we estimate has the following form:

$$Spread = \alpha + \beta_1 Period + \beta_2 Int + \beta_3 Int^2 + \beta_4 Value + \beta_5 Value^2 + \beta_6 Transp + \epsilon \qquad (1)$$

Each observation of the dependent variable corresponds to the value of the spread for a given cohort, market and interval.[6] *Period* is the number of the experimental market within a

[6]This means that the number of observations potentially equals 2 340, computed as 13 cohorts × 18 markets × 10 intervals. However, we do not have a valid spread value if there is no quote at the bid or at the ask, so the final number of observations equals 2 306.

session, that can take values from 1 to 18. It is expected to negatively influence the spread as a result of learning throughout the course of a session. Figure 2 shows that the spread decreases as trading progresses within a market, in a quadratic (convex) way. So we include both Int, one of the 10 intervals of a market expressed in seconds (15, 30, ..., 150), as well as its square value, Int^2. We also want to take into account the effect of the security value that has been randomly drawn. We expect the spread to be larger when the security has an "extreme" value, and to be lower when the random value is close to the expected value of the distribution – which is 50.5. That is why we include $Value$ and $Value^2$ as independent variables. Given the hypothesized convex relationships, we expect the coefficients on Int^2 and $Value^2$ to be positive. Finally, the $Transp$ variable is a dummy variable that takes the value 1 for markets played in a transparent setting and 0 otherwise. Hypothesis 1 dictates that the coefficient on $Transp$ should be negative, meaning that transparent markets have lower spreads and are thus more liquid.

Estimation results are presented in Table 1. In Panel A, we propose the results of equation (1), while Panel B displays the estimates for a slightly different specification, where we include 12 dummy variables to account for possible differences across cohorts. In Panel A, the intercept, Int and Int^2 are highly significant, while $Value$ and $Value^2$ have the expected sign but are marginally significant. Neither *Period* nor $Transp$ affects the bid-ask spread. However, the regression as a whole is not very good, with an R^2 of only 0.15. When we include the dummy variables accounting for differences across cohorts, things get better: the R^2 increases to 0.43, all variables are statistically significant with the expected sign, with the exception of the $Transp$ dummy variable that remains non significant.[7]

In the regressions above, many observations are provided by the same cohorts, and the independence hypothesis is not supported.[8] So we propose to analyze liquidity with an alternative method, i.e., a repeated-measures analysis of variance, the statistical procedure used by Bloomfield et al. (2005). Independence is ensured by keeping a single observation of the dependent variable for each cohort. Of course, the drawback of this method is that we lose much of the information. In our case, we perform an ANOVA with two within-subjects factors. *Extremity* measures the distance between a security true value and its prior expected

[7] We do not report the coefficients on the cohort dummies in order to conserve space.

[8] Considering that observations resulting from the same subjects are independent events may lead to an overstatement of statistical significance. This is probably not a major problem in our case, as the main variable of interest, $Transp$, is not significant. We also note that despite this caveat, regressions such as those we have performed are used in experimental studies – see for instance Flood et al. (1999) or Friedman (1993).

value. We classify markets in three extremity categories in the following way. We first compute the absolute difference (AD) between the security true value and the expected value ($AD \equiv |Value - 50.5|$). If $AD < 10$, then the extremity of a market is *low*; if $10 \leq AD < 30$, the market is of a *medium* extremity and if $AD \geq 30$ then the extremity of the market is *high*. The second factor is *transparency*, that takes two values: anonymous and transparent.

We first consider the value of the spread at the last interval. For each cohort, we compute the average value of the spread in each of the six cells that are defined by transparency and extremity, and the average across cohorts are provided in Table 2 (Panel A). There does not seem to be much variation due to one factor in particular, and this is confirmed by the ANOVA, that shows that there is no transparency main effect ($p = 0.6363$) and no extremity main effect ($p = 0.3298$). Results are qualitatively similar if we use the average spread over a market instead of the spread at the last interval (Panel B).

This means that anonymity in our experimental markets does not have any impact on the quoted spread, which does not support Hypothesis 1. Moreover, this is in contrast to the finding in Perotti & Rindi (2006) that anonymous markets are more liquid than transparent markets.

5.2 Informational efficiency

The second hypothesis deals with the relationship between anonymity and efficiency. As an inverse measure of efficiency, we use the absolute pricing error (APE), which is computed as the absolute difference between the quote midpoint and the true value of the security, i.e., $APE \equiv |Midpoint - Value|$. Figure 3 plots the evolution of APE over the trading game, with a distinction between transparency regimes. The result is consistent with the literature, as we can see a monotonic decrease of the absolute error over the trading intervals. This means that our experimental markets behave correctly as they progressively impound information into prices. As was the case for the bid-ask spread, no difference appears between the anonymous and transparent regimes.

We propose to use the same econometric specification as in Equation (1), using APE as the dependent variable. Results are provided in Table 3. In both specifications – with and without cohort dummies – all control variables are highly significant with the expected sign, and the $Transp$ dummy is not significant at all. As in the previous section, we also use an ANOVA procedure to analyze the absolute pricing error. The average values across cohorts of final APE by transparency and extremity are shown in Panel A of Table 4. If APE seems to increase with extremity, the difference is not statistically significant ($p = 0.4148$). There is also no effect due to transparency ($p = 0.5600$). If we consider the average APE over a trading game rather than the final value (see panel B of Table 4), then the extremity factor becomes statistically significant ($p = 0.0006$) but transparency still has no impact.

In summary, displaying the ID codes of traders in the order book does not reduce the pricing error and so does not make the market more efficient. We do not find any support for Hypothesis 2, but this is in line with the results provided by Perotti & Rindi's (2006) experimental markets. This result is however inconsistent with the traditional view that a higher transparency level improves market efficiency.

5.3 Traders' profits

The third hypothesis states that informed traders should be better off in an anonymous setting, while liquidity traders should benefit from transparency. So we compute the end-of-period profits for the three trader categories – informed traders, buyers and sellers. More precisely, we compute *trading profits*, i.e., we do not take into account the possible penalties incurred by liquidity traders. In our markets, 82.05% of buyers (78.63% of sellers) reach their objective, and these figures are not affected by the transparency regime. For those liquidity traders who do not meet their target, the median penalty incurred equals 40 ECU, and this figure is the same for buyers and sellers and in both transparency regimes. This means that those traders do not reach their objective by only 2 shares. However, it is not by lack of time. Indeed, buyers and sellers who reach their objective do it after an average of 83 seconds, with once again no difference between transparency regimes. Given that a market lasts 150 seconds, traders have plenty of time to fulfill their target. Note also that there is no learning effect in that matter, as the proportion of traders who do not meet their target is very stable across the 18 periods of a session.

Table 5 provides the average values of trading profits for the three trader types, by transparency regime and extremity. We expect the profits of the informed traders to be higher when the value of their information is high, i.e., when the true security value is far from its prior expected value. This is indeed what we find: profits for the informed traders monotonically increase with extremity, in both transparency regimes. Performing an ANOVA on those data shows that the extremity variable is highly significant ($p = 0.0011$). Informed traders seem to earn less profits in the anonymous regime with low extremity and more profits when the market is anonymous and extremity is high. However, transparency is not significant in the

ANOVA results ($p = 0.2312$), and there is also no transparency $\times$ extremity interaction effect ($p = 0.9929$).

Results for buyers and sellers call for several comments. Firstly, those traders logically incur losses, but those losses are lower than the penalty they would incur if they decided not to trade at all – remember that a liquidity trader who does nothing has a penalty of $5 \times 20 = 100$ ECU. Liquidity traders are thus better off trading than not trading. In the same way as informed traders earn more profits when extremity is high, buyers have bigger losses when the security value is far from its expected value ($p = 0.0030$ for extremity). However this is not true for sellers, as their losses in the transparent/high extremity cell are lower in absolute value than in the transparent/medium extremity cell. The p-value associated to the extremity factor is at the margin of significance, as it equals 0.1202. As for transparency, it does not affect the profits of buyers ($p = 0.7093$), but it is significant for sellers at the 10% level ($p = 0.0679$). However, if we perform a test of difference (t-test or Wilcoxon test) between the anonymous and transparent regimes in each of the three extremity classes for sellers, none of them is statistically significant. Also, if we make the same analysis where we group buyers and sellers into a single "liquidity traders" group, then extremity is highly significant while transparency is not.

The asymmetry in the impact of extremity on buyers' and sellers' profits is puzzling at first sight. It seems however that it is the transparent/high extremity cell that causes the problem. This may find an explanation. Indeed, the high extremity class contains very low security values (under 20) as well as very high values (above 80). Intuitively, buyers should be in a better position when the true value is high, while the opposite should hold for sellers. This is confirmed by figures provided in Table 6, where we present the average profits of buyers and sellers in those extreme cases. We see that sellers are even able to make profits in transparent markets when the value is extremely low – but the difference between anonymous and transparent regimes is not statistically significant, as is the case of all four cells of that table.

It turns out that more markets have been characterized by an extremely low value (28 in the anonymous regime and 24 in the transparent regime) than by an extremely high value (16 and 17 in the anonymous and transparent cases respectively). This is due to the random generating process of the security true value, and partly explains why losses for sellers in the high extremity cell of Table 5 are low (in absolute value), and in particular lower than for buyers.

In summary, we do not find any evidence in support of Hypothesis 3, i.e., transparency does not influence the welfare of traders of any kind. Note that Perotti & Rindi (2006) do not analyze traders' profits.

5.4 Traders' behavior

The three previous sections have shown that anonymity in our experimental markets does not have any influence on liquidity, efficiency or traders' profits, thus providing support to none of the hypotheses developed in Section 3. In this section, we look more closely at the behavior of traders in the experimental markets, in order to understand why the knowledge of informed traders' identity does not play any role. Given that there is no difference between both transparency regimes, from now on we present results that are aggregated across all markets.

We first focus on the different types of orders that are placed by traders. In our experiment, traders can submit two types of orders: market orders and limit orders. Limit orders are further disaggregated in three categories: limit orders that are executed, limit orders that traders actively cancel and limit orders that are left unexecuted in the order book at the end of the trading period. Figure 4 provides for each trader type the average number of orders of those four categories that are introduced during a market.

The behavior of buyers and sellers is very similar, and is different from the behavior of informed traders. Informed traders place more orders than liquidity traders (21 against 16), and both more market orders (7 against 5) and more limit orders (14 against 11). The sum of market orders and executed limit orders equals the number of transactions, which is slightly higher for informed traders (13) than for liquidity traders (11). Note that this means that approximately 52 shares change hands during a typical market. There are very few cancellations. For buyers and sellers, 52% of their limit orders are executed, but that figure drops to 41% for informed traders.

This graph is very similar to Bloomfield et al.'s (2005) results, thus showing that traders in our experimental markets behave in the same way as traders in their experiment. The picture that emerges is one where informed traders act not only as liquidity takers, but also as important liquidity suppliers. This is consistent with their behavior in theoretical models developed by Kaniel & Liu (2006), Moinas (2006) and Rindi (2004), and with the empirical findings in Kaniel & Liu (2006) that limit orders are more informative than market orders.

In order to better understand the dynamics of our markets, we now look more closely at how traders' behavior evolves during a trading game. So we compute the *taking rate*, a measure that has been proposed by Bloomfield et al. (2005), defined as the percentage of trades completed using market orders (i.e., the number of market orders divided by the sum of the market orders and the executed limit orders). The evolution of the taking rate over a game for the three categories of players is displayed on Figure 5.

For sellers, the taking rate starts at 41%, the lowest level among all types, remains stable for four intervals, then increases toward 46 to 48%. This means that sellers are relatively patient at the beginning of the game, and then become increasingly aggressive as trading progresses. This is in sharp contrast to the behavior of buyers, who are most aggressive at the start of the game, with a taking rate of 52%, that decreases progressively until the seventh interval, where it suddenly increases to reach the same level as is observed for sellers. This aggressive behavior of buyers may explain why they have bigger losses than sellers, as they act as liquidity takers at the beginning of trading, i.e., when quotes are still far from the security true value.

A third pattern characterizes the evolution of the taking rate for informed traders. It starts at 45%, the lowest level observed for this trader type, but it quickly increases to reach a peak at 52% during the third interval. Then it slowly decreases troughout the trading game, but it remains higher than 50% for four intervals, and also shows an increase in the next-to-last interval. This pattern shows that informed traders take profit opportunities at the start of the game, by hitting quotes that are mispriced – bid quotes above the true value and ask quotes under that value. Two elements probably account for the sharp increase that can be observed in the second and third intervals: (i) the taking rate cannot be high at the very start of the game, given that the order book is empty, so that informed traders have to wait until some profit opportunities appear and (ii) there is a competition between informed traders to take those profits as soon as they appear.

This last hypothesis is confirmed by looking at the evolution of the average profit by transaction for each type of trader, which is provided in Figure 6. This graph clearly shows that the best profit opportunities appear at the beginning of the game. The average profit by transaction for informed traders is indeed highest at that time, and then monotonically decreases as trading progresses. The opposite pattern logically holds for buyers and sellers, and we may notice that the losses of buyers for the first five intervals are greater than the losses of sellers, which is consistent with the evidence presented above that buyers are more aggressive at the beginning of a market.

We now provide some details concerning the provision of liquidity. We adapt the concept of order aggressiveness defined by Biais et al. (1995), and classify each order in one of four aggressiveness categories: (i) the first category contains market orders, the only orders in our experimental setting leading to an immediate transaction; (ii) the second category is made of limit orders that improve on the existing best quote; (iii) in the third class, we find limit orders that add depth at the best quote; (iv) orders with a price below the best quote form the last category. Aggressiveness is defined in such a way that it decreases from the first to the fourth category. Figure 7 provides for each trader type the evolution of the proportion of the different order categories.

We first note that the evolution for buyers and sellers is quite similar. The proportion of orders of categories 3 and 4 is relatively stable over the trading game, between 15 and 20%. At the beginning of the game, sellers tend to place many orders of the second category, i.e., orders that improve on the best quote, but after 60 seconds there is a switch to market orders, that become the most used category. The pattern is a little less pronounced for buyers. Informed traders always submit more market orders, except in the first interval for the reason we have mentioned above. In terms of provision of liquidity, their behavior is also stable over time. In contrast to liquidity traders, they submit more orders of the third category than of the second category. In other words, rather than improving on the best quote, they add depth at the existing quote.

To provide more insights in this matter, Figure 8 displays the contribution of the three types of traders to depth at the best bid and ask limits. More precisely, it plots for both market sides the depth available at the best limit, disaggregated according to trader type.

We first highlight that depth is bigger on the ask side than on the sell side. On the bid side, buyers contribute the most to depth for the first five intervals, and are then replaced by informed traders who become the best depth contributors and remain so until the end of trading. Sellers logically contribute less to depth at the best bid. The picture for the ask side is a little different, as informed traders are always the best liquidity providers – but they tie with sellers for 60 seconds –, followed by sellers and then buyers. It can be surprising that informed traders are the best liquidity providers on the ask side since the beginning of the trading game. One possible explanation is that ask quotes are farther away from the true value, so that informed traders do not reveal much of their information by being at the best limit at the beginning of the game. To confirm this hypothesis, we compute, for each limit order that is introduced, the difference between the price and the true security value. For the first interval of the game, the average value of this difference in the case of informed traders equals 17 for buy orders and 27 for sell orders. The corresponding values for the second interval are 11 (buy orders) and 17 (sell orders). This means that, consistent with our hypothesis, informed traders post ask quotes that are conservative and do not reveal much information to the market.

Our analysis of traders' behavior helps explain why anonymity does not play a role in those markets. This is due to the way informed traders behave. At the beginning of the game, they mainly act as liquidity takers, as they are in competition with each other in order to reap as much profits as they can. As trading progresses, they become real liquidity providers, in the sense that they provide liquidity at the best quote. This is the moment where the disclosure of ID codes could play a role, but it is actually too late: the information has been revealed to the market, the true price has been discovered, and the main profits have been earned.

6 Increasing the proportion of informed traders

The results presented in the previous section have shown that anonymity does not influence traders' behavior or market quality in our experimental markets. This is in contrast to the results provided by Perotti & Rindi (2006), who show that anonymity, by inducing more traders to become informed, improves market liquidity. This emphasizes the important role played by the market for information in their design. Indeed, when studying liquidity, Perotti & Rindi establish a link between the spread decrease and the increase in the number of informed traders. However this does not mean that the incremental information that is displayed during the market game – the ID codes of traders in the order book – is taken into account by players.

In this section we propose to *exogenously* increase the proportion of informed traders in our experimental markets. The design we implement is exactly the same as the one we have used so far, except for the number of players in the different roles. While we keep nine subjects who participate in each session, five of them are now informed (instead of three in the previous design), two act as buyers and the last two as sellers (there were three buyers and three sellers in the original design).[9]

If Perotti & Rindi's (2006) results are mainly due to the proportion of informed traders, then we expect to observe the same effects, i.e., we expect that markets in the second design will be more liquid than markets in the first design. This finding would also be consistent with the theoretical predictions of Holden & Subrahmanyam (1992). These authors extend the Kyle

[9]The proportion of informed traders in our first and second designs tries to match the proportion of informed traders in the transparent and anonymous markets of Perotti & Rindi (2006). In transparent markets, 50% of their 12 players on average buy the information. Adding the six robots, this means that 6 players out of 18 are informed, which equals the proportion of informed traders in our first design (3 out of 9). In the anonymous setting, Perotti & Rindi (2006) indicate that 71.5% of traders decide to become informed, or 8.6 subjects on average. This gives a proportion of 8.6/18, that is 47.8% of informed traders, while we have a little more in our second design: 55.6% (5 subjects out of 9).

(1985) model, by having multiple privately informed agents instead of a monopolistic trader. In that setting, an increase in the number of informed traders improves both market liquidity and efficiency, because aggressive competition between those traders causes their information to be revealed very quickly, and market depth to become extremely large. This means that we also expect markets in the second design to be more informationally efficient than in the first design, and also more volatile.

Seven sessions have been conducted under this second design, with other unexperienced students at the Catholic University of Mons. Figure 9 displays the evolution of the quoted spread, the absolute pricing error and volatility (computed as the standard deviation of transaction prices) in the anonymous and transparent structures of both experimental designs. If we look only at the results of the second design (the "D2" curves), we notice that, as was the case with the first design, no difference appears between both transparency regimes. We have performed the same statistical tests as in the previous section (multivariate regressions and repeated-measures ANOVA), and all of them point to the absence of an effect due to the transparency regime.[10]

What is striking however is that *all* the curves pertaining to that second design lie under the corresponding curves of the first design. In other words, using a higher proportion of informed traders leads to better liquidity, to more informationally efficient and less volatile markets. The improvement is really spectacular: for instance, the average absolute pricing error in the second design is no more than 3 ECU after only 60 seconds. Even if the difference between both designs is very clear by looking at the graphs only, we also address it in a multivariate framework. For the three market quality measures, we use the specification of Equation (1), where we add another explanatory variable, $Design2$, a dummy that equals 1 for sessions conducted under the second experimental design and 0 otherwise. The results are given in Table 7.

The econometric results clearly demonstrate that the markets played under the second design are more liquid and efficient, and also less volatile. The effect is economically significant, as the average reduction in spread equals 8.55 ECU, the average decrease in pricing error 8.69

[10]The second graph of Figure 9 may give the impression that the pricing error of the second design is larger in the anonymous markets than in the transparent markets during the first 6 intervals. To investigate whether this is the case, we have performed two analyses. Firstly we have added an interaction variable between Int and $Transp$ in specification (1), that would account for a possibly differing slope of the Int variable between transparency regimes. However, it turns out that this new variable is not significant. We have also estimated equation (1) by using only data from the first six intervals, and the $Transp$ variable remains unsignificant.

ECU and the average reduction in volatility 3.79 ECU. Results on liquidity and efficiency are consistent with Holden & Subrahmanyam's (1992) theoretical model, predicting that both quality measures should be improved thanks to the competition between informed traders. In our experimental setting, this competition drives down their profits. Indeed, the average profits of informed traders equal 10.49 in the low extremity class, 8.37 in the medium extremity and 15.88 in the high extremity. All these figures are lower than the profits earned by informed traders in the first experimental design (see Table 5).

The results on volatility are more surprising, because the usual hypothesis is that informed traders induce more volatile markets. Perotti & Rindi (2006) expect their anonymous markets to display more volatility given the highest number of informed traders.[11] While transparent markets in their experiment are indeed less volatile, the authors are not able to make a link with the number of informed traders.

Our finding that liquidity is improved by an exogenous increase in the number of informed traders helps to shed a new light on Perotti & Rindi's (2006) results. In particular, if anonymity plays a role in their markets, it is very likely that it only comes from the impact due to the market for information and not from the varying information that is displayed during the trading game.

As mentioned in Section 2, Sunder (1992) analyzes how the organization of a market for information affects the equilibrium that prevails in the subsequent trading market. When information is sold to a fixed number of highest bidders, the predictions of the full rational expectations equilibrium (REE) can be observed in the trading: the asset prices and asset allocation among participants converge to REE predictions, and the price of information converges to zero. This last finding reflects that demand for information shifts to the left, consistent with the idea that information is useless as prices reveal all information. But when the information is sold at a fixed price to a varying number of traders, as is the case in Perotti & Rindi (2006), then markets are less informative and the number of traders who buy the information does not converge to zero. Sunder's (1992) findings suggest that other results could have been obtained by Perotti & Rindi (2006) if they had decided to let players bid for the information, instead of paying a fixed price. They may also provide an explanation for why Perotti & Rindi (2006) do not find any statistically significant effect on market efficiency.

[11] Note that while transparency affects volatility in Rindi's (2004) model, the effect depends on the parameters under consideration.

7 Conclusion

In this paper, we have performed an analysis of the impact of pre-trade anonymity within the context of an experimental market. More precisely, we have implemented an electronic double-auction market, similar in many ways to the order-driven system that is used on many exchanges around the world, and we have compared the outcomes of two settings differing only in the degree of information disclosed to market participants: in a transparent setting, traders see the ID codes alongside the orders in the order book and know who the informed traders are, while the ID codes are not disclosed in an anonymous structure.

It turns out that displaying or not the ID codes of traders does not affect traders' behavior, and as a corollary does not have any influence on market quality indicators – liquidity, efficiency and volatility. This result may be explained by the fact that informed traders are in competition with each other in order to get the highest profits, and so behave aggressively at the beginning of the game by taking liquidity, thus revealing their information. When they start to act as real liquidity providers, i.e., when the information relative to their ID code may be helpful, the game is over.

Our results are in sharp contrast to the findings of Perotti & Rindi (2006), who experimentally show that anonymity improves market liquidity. In their experiments, the actual trading game is preceded by a market for information, during which traders can choose whether they want to buy a signal on the liquidation value of the asset. In that design, anonymity induces more traders to buy the information, which leads the subsequent markets to be more liquid, because informed traders are the best liquidity providers. This is consistent with the endogenous information acquisition version of the theoretical model developed by Rindi (2004). Our results however do not support the theoretical predictions developed by Rindi (2004) in the version with a constant number of informed traders.

This means that the presence or absence of a market for information is really key in deriving the possible implications of a change in anonymity on market quality. If traders indeed change their behavior in terms of information gathering, then it may be the case that a switch to anonymity improves market quality. However, if the balance between uninformed and informed traders is not affected by anonymity, then we should not observe any change.

Finally, it can be interesting to compare our "no-impact" result with results from other papers in the literature focusing on the effects of pre-trade transparency in experimental markets. Flood et al. (1999) use a continuous quote-driven market with multiple dealers, and show that quote disclosure does affect market liquidity and efficiency. However, there is an important gap in terms of information disclosure between the structures they compare. In their “opaque” market, no quote is publicly disclosed and market makers have to call each other in order to obtain quote information, while all market makers’ quotes are disclosed in their “transparent” setting. Friedman (1993) analyzes how the order book disclosure affects the performance of a continuous double auction. He shows that displaying *all* the bid and ask quotes instead of the *best limit* only seems to improve market efficiency, possibly at the expense of liquidity. In this case also, the difference in the amount of information available to traders was relatively important.

By contrast, the transparency regimes we have compared in this paper differ only by one variable, the ID code of traders. In other words, the anonymous markets in our experiments are already very transparent as they disclose all quotes to all participants. That difference between both structures mimics the difference between the trading platforms prevailing on Euronext Paris or the Tokyo Stock Exchange before and after they switched to anonymity. Our findings seem to show that the incremental information contained in the identity of traders does not suffice to change the market equilibrium. This conclusion is consistent with the conjecture made by Plott & Sunder (1982) that bids and offers, more than the ability to identify insiders, convey most of the information.

References

Admati, A. R. & Pfleiderer, P. (1991), ‘Sunshine trading and financial market equilibrium’, *The Review of Financial Studies* **4**(3), 443–481.

Albanesi, S. & Rindi, B. (2000), ‘The Quality of the Italian Treasury Bond Market, Asymmetric Information and Transaction Costs’, *Annales d’Economie et de Statistique* **60**, 1–19.

Barclay, M. J., Hendershott, T. & McCormick, D. T. (2003), ‘Competition among trading venues: Information and trading on electronic communications networks’, *The Journal of Finance* **58**(6), 2637–2665.

Battalio, R., Ellul, A. & Jennings, R. (2007), ‘Reputation Effects in Trading on the New York Stock Exchange’, *The Journal of Finance* **62**(3), 1243–1271.

Benveniste, L. M., Marcus, A. J. & Wilhelm, W. J. (1992), ‘What’s special About the Specialist?’, *Journal of Financial Economics* **32**(1), 61–86.

Bessembinder, H. & Venkataraman, K. (2004), 'Does an Electronic Stock Exchange Need and Upstairs Market?', *Journal of Financial Economics* **73**(1), 3–36.

Biais, B., Hillion, P. & Spatt, C. (1995), 'An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse', *The Journal of Finance* **50**(5), 1655–1689.

Bloomfield, R. & O'Hara, M. (1999), 'Market Transparency: Who Wins and Who Loses?', *The Review of Financial Studies* **12**(1), 5–35.

Bloomfield, R., O'Hara, M. & Saar, G. (2005), 'The "Make or Take" Decision in an Electronic Market: Evidence on the Evolution of Liquidity', *Journal of Financial Economics* **75**(1), 165–199.

Bloomfield, R., O'Hara, M. & Saar, G. (2006), 'The Limits of Noise Trading: An Experimental Analysis', *Working Paper.* Johnson Graduate School of Management, Cornell University.

Boehmer, E., Saar, G. & Yu, L. (2005), 'Lifting the Veil: An Analysis of Pre-Trade Transparency at the NYSE', *The Journal of Finance* **60**(2), 783–815.

Booth, G. G., Lin, J.-C., Martikainen, T. & Tse, Y. (2002), 'Trading and pricing in upstairs and downstairs stock markets', *The Review of Financial Studies* **15**(4), 1111–1135.

Comerton-Forde, C., Frino, A. & Mollica, V. (2005), 'The impact of limit order anonymity on liquidity: Evidence from Paris, Tokyo and Korea', *Journal of Economics and Business* **57**(6), 528–540.

Desgranges, G. & Foucault, T. (2005), 'Reputation-based pricing and price improvements', *Journal of Economics and Business* **57**(6), 493–527.

Dia, M. & Pouget, S. (2006), 'Sunshine trading in an emerging stock market', *Working Paper.*

Easley, D. & O'Hara, M. (1987), 'Price, trade size, and information in securities markets', *Journal of Financial Economics* **19**(1), 69–90.

Eom, K. S., Hahn, S. B. & Seon, J. (2006), 'Traders' anonymity, market quality, and foreign traders' superior information: Evidence from the Korea Exchange', *Working Paper.*

Fischbacher, U. (2007), 'z-tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics* **10**(2), 171–178.

Fishman, M. J. & Hagerty, K. M. (1995), 'The mandatory disclosure of trades and market liquidity', *The Review of Financial Studies* **8**(3), 637–676.

Fishman, M. J. & Longstaff, F. A. (1992), 'Dual trading in futures markets', *The Journal of Finance* **47**(2), 643–671.

Flood, M. D., Huisman, R., Koedijk, K. G. & Mahieu, R. J. (1999), 'Quote disclosure and price discovery in multiple-dealer financial markets', *The Review of Financial Studies* **12**(1), 37–59.

Forster, M. M. & George, T. J. (1992), 'Anonymity in securities markets', *Journal of Financial Intermediation* **2**(2), 168–206.

Foucault, T., Moinas, S. & Theissen, E. (2007), 'Does anonymity matter in electronic limit order markets?', *The Review of Financial Studies* **20**(5), 1707–1747.

Friedman, D. (1993), 'How trading institutions affect financial market performance: Some laboratory evidence', *Economic Inquiry* **31**(3), 410–435.

Frino, A., Johnstone, D. & Zheng, H. (2005), 'The information content of trader identification', *Working Paper*. University of Sydney.

Garfinkel, J. A. & Nimalendran, M. (2003), 'Market structure and trader anonymity: An analysis of insider trading', *Journal of Financial & Quantitative Analysis* **38**(3), 591–610.

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

Grammig, J., Schiereck, D. & Theissen, E. (2001), 'Knowing me, Knowing you: Trader Anonymity and Informed Trading in Parallel Markets', *Journal of Financial Markets* **4**(4), 385–412.

Grossman, S. J. & Stiglitz, J. E. (1980), 'On the impossibility of informationally efficient markets', *The American Economic Review* **70**(3), 393–408.

Heidle, H. G. & Huang, R. D. (2002), 'Information-based trading in dealer and auction markets: An analysis of exchange listings', *Journal of Financial & Quantitative Analysis* **37**(3), 391–424.

Holden, C. W. & Subrahmanyam, A. (1992), 'Long-lived private information and imperfect competition', *The Journal of Finance* **47**(1), 247–270.

Jain, P. K., Jiang, C., McInish, T. H. & Taechapiroontong, N. (2006), 'Informed trading in parallel auction and dealer markets: An analysis on the London Stock Exchange', *Working Paper*.

Kaniel, R. & Liu, H. (2006), 'So what orders do informed traders use?', *Journal of Business* **79**(4), 1867–1913.

Kyle, A. S. (1985), 'Continuous auctions and insider trading', *Econometrica* **53**(6), 1315–1335.

Madhavan, A. (1996), 'Security prices and market transparency', *Journal of Financial Intermediation* **5**(3), 255–283.

Madhavan, A. & Cheng, M. (1997), 'In search of liquidity: Block trades in the upstairs and downstairs markets', *The Review of Financial Studies* **10**(1), 175–203.

Madhavan, A., Porter, D. C. & Weaver, D. G. (2005), 'Should securities markets be transparent?', *Journal of Financial Markets* **8**(3), 266–288.

Majois, C. (2007), 'Natural experiments methodology and global liquidity in financial markets', *Working Paper*.

Moinas, S. (2006), 'Hidden limit orders and liquidity in limit order markets', *Working Paper*. Toulouse Business School.

O'Hara, M. (1995), *Market Microstructure Theory*, Blackwell.

Perotti, P. & Rindi, B. (2006), 'Market for information and identity disclosure in an experimental open limit order book', *Economic Notes* **35**(1), 97–119.

Plott, C. R. & Sunder, S. (1982), 'Efficiency of experimental security markets with insider information: An application of rational-expectations models', *Journal of Political Economy* **90**(4), 663–698.

Reiss, P. C. & Werner, I. M. (2005), 'Anonymity, adverse selection, and the sorting of interdealer trades', *The Review of Financial Studies* **18**(2), 599–636.

Rindi, B. (2004), 'Informed traders as liquidity providers: Transparency, liquidity and price formation', *Working Paper*. Bocconi University.

Roëll, A. (1990), 'Dual-capacity trading and the quality of the market', *Journal of Financial Intermediation* **1**(2), 105–124.

Scalia, A. & Vacca, V. (1999), 'Does market transparency matter? A case study', *CGFS Publications* **11**.

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

Simaan, Y., Weaver, D. G. & Whitcomb, D. K. (2003), 'Market maker quotation behavior and pretrade transparency', *The Journal of Finance* **58**(3), 1247–1267.

Smith, B. F., Turnbull, D. A. F. & White, R. W. (2001), 'Upstairs market for principal and agency trades: Analysis of adverse information and price effects', *The Journal of Finance* **56**(5), 1723–1746.

Sunder, S. (1992), 'Market for information: Experimental evidence', *Econometrica* **60**(3), 667–695.

Theissen, E. (2002), 'Floor versus screen trading: Evidence from the German stock market', *Journal of Institutional and Theoretical Economics* **158**(1), 32–54.

Theissen, E. (2003), 'Trader anonymity, price formation and liquidity', *European Finance Review* **7**(1), 1–26.

Waisburd, A. C. (2004), 'Trade transparency, identity disclosure, and the cost of transacting', *Working Paper*.

A Instructions handed to participants

You are about to participate to an experiment on financial markets. If you understand the instructions correctly and make the right decisions, you can win a consequential sum of money.

During the experiment, we will use a fictive currency: the ECU (Experiment Currency Unit). Your final wealth in ECU at the end of the experiment will be converted into euros.

Communication between subjects is forbidden during the course of the experiment. If one subject tries to communicate with another one, or tries to look at another subject's computer, he will be excluded from the experiment without any payment.

Please carefully read the instructions hereunder. In these instructions, the terms in italic refer to information or actions you will see on the computer screens. When you will have read instructions, we will review them together and you will have the opportunity to ask questions.

The market

During the experiment, you will play several market games.[12] During each game, you will have the opportunity to buy or sell several units of a security. At the beginning of each game, the value of this security is drawn randomly and independently of the other games. This value may be any integer number between 1 and 100 ECU, each with the same probability.

You will trade this security in an order-driven market. In this kind of market, the displayed order book contains buy and sell orders waiting for execution. Here is a sample order book:

Bid prices	Ask prices
78	81
78	87
45	87
	98

The left column (*Bid prices*) contains the buy orders waiting for execution. These orders are ranked from the highest price (at the top of the book) to the lowest (at the bottom). The right column (*Ask prices*) contains the sell orders waiting for execution. These orders are ranked from the lowest price to the highest. This means that the best buying and selling opportunities are found at the top of the book. In the above example, the best price proposed by a buyer is 78, and the best price proposed by a seller is 81. Those prices represent the "best

[12] Note that the expressions "game" and "period" have the same meaning in the following.

limit" of the order book. If several orders inside a column have the same limit price, they are ranked depending on the time of their introduction: the order that has been introduced first is above the others (there is a time priority).

If you want to buy one unit of the security, you may choose between 2 options:

- Buy immediately (*Buy*). You buy one unit of the security at the price proposed by the best seller, i.e., at the price displayed at the best ask limit in the order book (81 in the above example). In that case, the corresponding sell order disappears from the book.
- Make a bid proposal, by introducing a buy order with a limit price (*Propose Bid*). This order appears in the *Bid prices* column. There is no immediate transaction, but only a trade proposal. There will be a trade if your order is accepted by another player.

Similarly, if you want to sell one unit of the security, you can:

- Sell immediately (*Sell*). You sell one unit of the security at the price proposed by the best buyer, i.e., at the price displayed at the best bid limit in the order book (78 in the above example). In that case, the corresponding buy order disappears from the book.
- Make an ask proposal, by introducing a sell order with a limit price (*Propose Ask*). This order appears in the *Ask prices* column. There is no immediate transaction, but only a trade proposal. There will be a trade if your order is accepted by another player.

You may also cancel (*Cancel*) a buy or sell proposal that you have made earlier in the game.

All the orders are for one unit of the security only. If you want to trade more, you may place several orders at the same price.

You have lots of freedom in the actions you can take. However, you cannot:

- trade with yourself;
- place a buy (sell) order with a price equal to or higher (lower) than the best ask (bid) price in the order book;
- buy or sell at a price different from the prices at the best limit;
- cancel an order placed by another player.

Error messages will be displayed on the screens if you try to take one of those actions.

Note that there is no obligation to trade.

Trader types

There are three types of players in each game:

- Players with a “buy objective”: these players must complete at least 5 buy transactions during the game. For every trade that is not completed, a 20 ECU penalty will be deducted from the final wealth of the player.
- Players with a “sell objective”: these players must complete at least 5 sell transactions during the game. For every trade that is not completed, a 20 ECU penalty will be deducted from the final wealth of the player.
- Players without any specific objective, but who know the security value that has been randomly drawn.

At the beginning of each game, player roles are randomly assigned to subjects. Each subject will have the opportunity to play all roles during the experiment.

How earnings are computed

Each player starts each game with 0 unit of security and 0 ECU in cash. There is no negative constraint about the amount of security or ECU that you may have: you may buy securities even if you have no money (there is no “borrowing constraint”), and you may sell a unit of the security even if you do not have it in stock (there is no “short selling constraint”).

Your final wealth at the end of a game is computed as follows:

$$\text{Cash amount} + \text{Value of your security portfolio} - \text{Penalty}$$

<u>Example 1</u>

You have been assigned a buy objective. During the game, you have made 2 buy trades at 15 and 25 ECU, and you have sold one unit at 50 ECU. At the end of the game, you have made only 2 buy transactions instead of 5.

Your cash amount is 10 ECU (50 − 15 − 25), and you have one security unit in your portfolio. The randomly drawn true value for the security was 35 ECU. Your earnings are:

$$10 \text{ (Cash)} + 1 \text{ X } 35 \text{ (Portfolio)} - 3 \text{ X } 20 \text{ (Penalty)} = -15 \text{ ECU}$$

Example 2

You had no objective. You have bought 5 units at respectively 45, 67, 70, 72 et 80 ECU, and sold 1 unit at 88 ECU.

At the end of the game, you have − 246 ECU in cash (88 − 45 − 67 − 70 − 72 − 80) and 4 units in portfolio. The true value was 72 ECU. Your final wealth is:

$$- 246 \text{ (Cash)} + 4 \text{ X } 72 \text{ (Portfolio)} = 42 \text{ ECU}$$

You final wealth in ECU at the end of the experiment is computed as the sum of your final wealth at the end of each game.

This final wealth in ECU will be converted into euros at the end of the experience, in such a way that the minimum payment will be 10 euros and the average payment 20 euros. Note that a negative wealth in ECU does not mean that you will get the minimum payment.

Market structures

You will play in two different market structures. The only difference between the structures relates to the information that is displayed to you:

- In an "anonymous" market, you will see order prices in the order book.
- In a "transparent" market, you will see order prices, the ID codes of players who placed those orders, and you will know who are the informed players.

Number of games and duration

You will play 22 games during this experiment. Each game lasts 150 seconds.

The first four games are practice games. These games are played in order for you to familiarize yourself with the market system and the computer screens. The earnings you will get during those four practice games will not be taken into account for the computation of your final wealth at the end of the experiment.

The earnings you will make during the remaining 18 games will be used as described earlier to compute your earnings in euros for this experiment.

Figure 1: Trading screen in a transparent regime

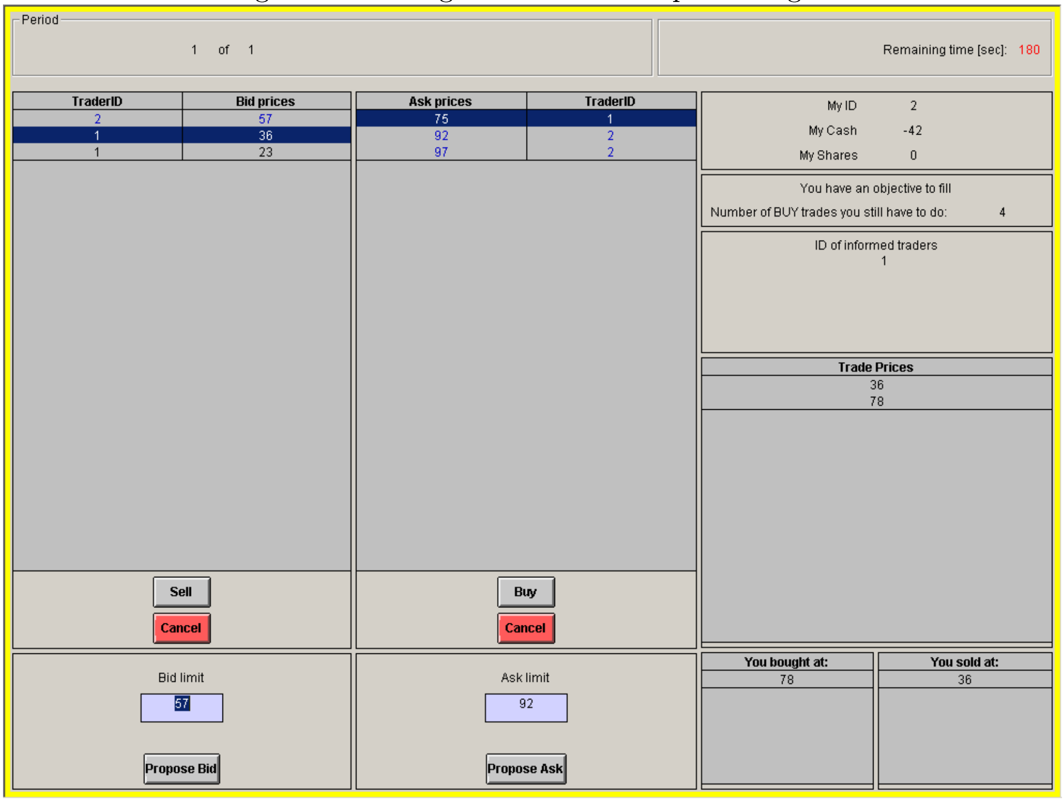

Figure 2: Quoted spread over the trading intervals

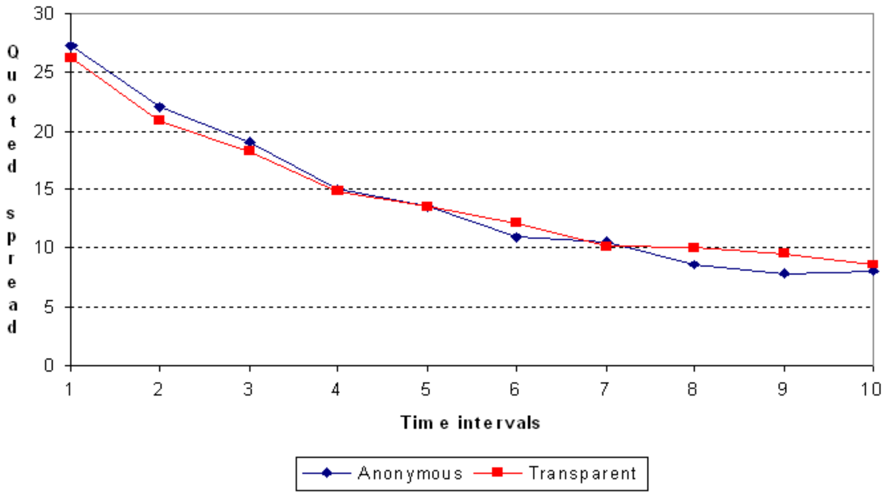

Figure 3: Absolute pricing error over the trading intervals

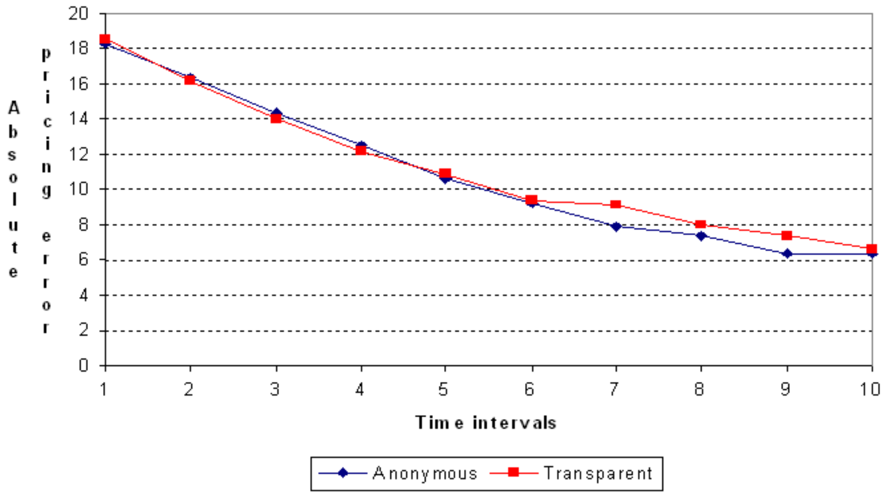

Figure 4: Types of orders by player type

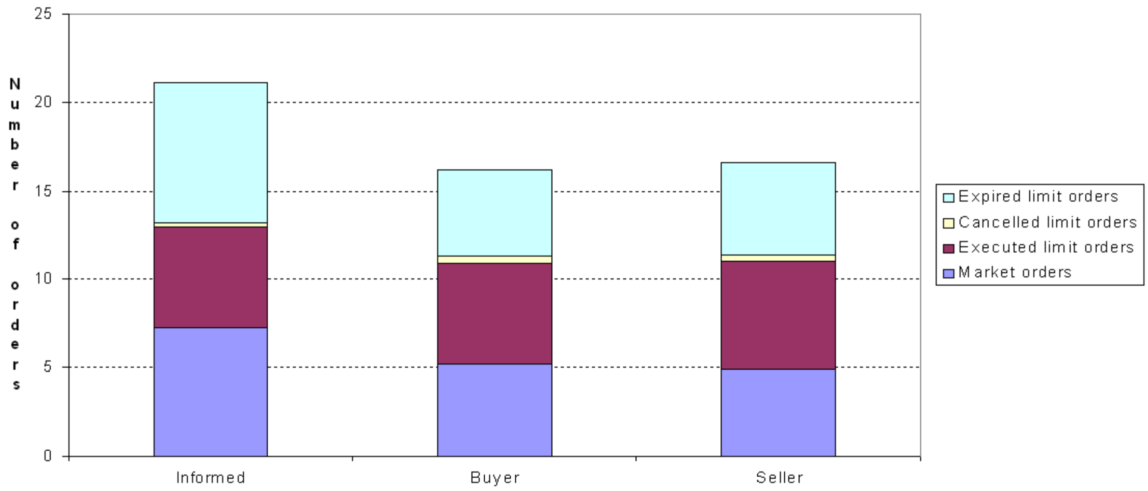

Figure 5: Taking rate over a game by player type

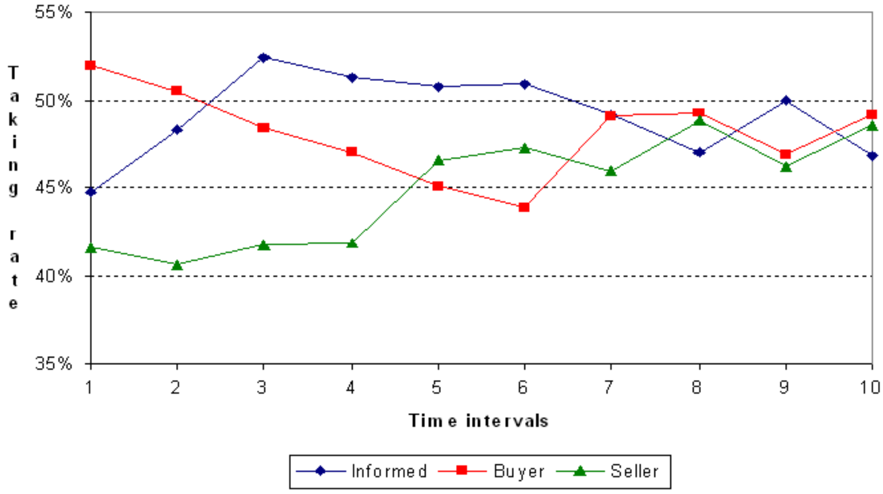

Figure 6: Evolution of the average profit by transaction

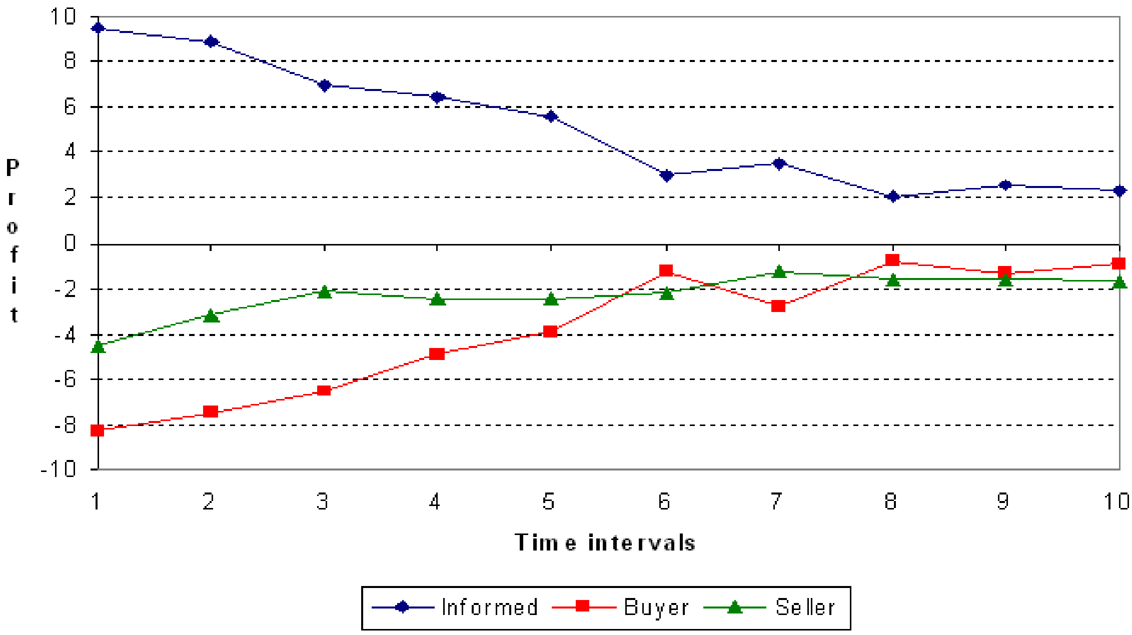

Figure 7: Order aggressiveness over the trading game

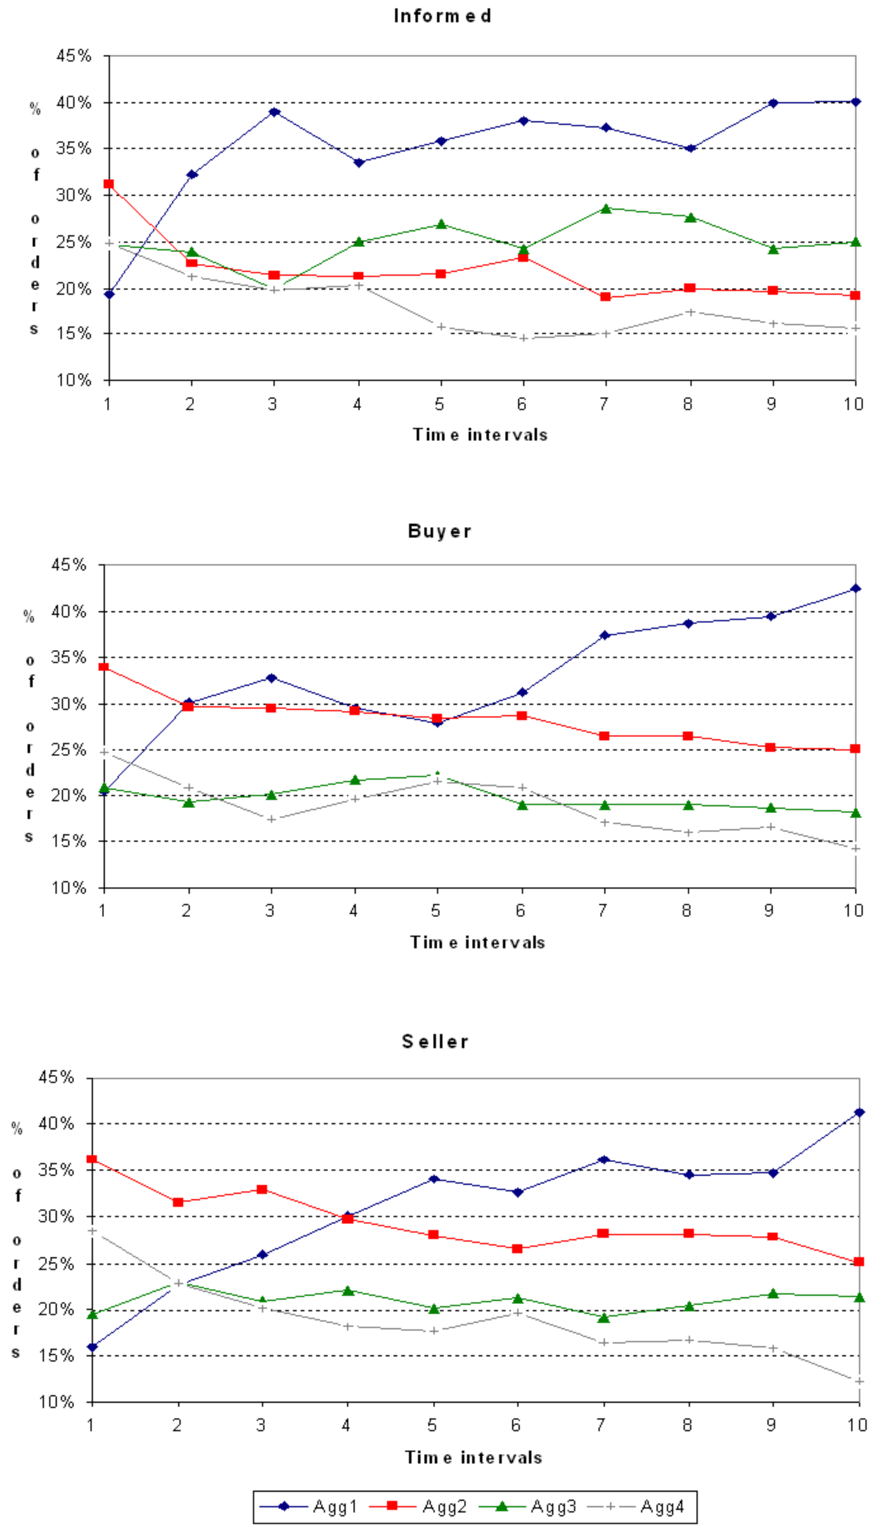

Figure 8: Contribution to depth at best limits

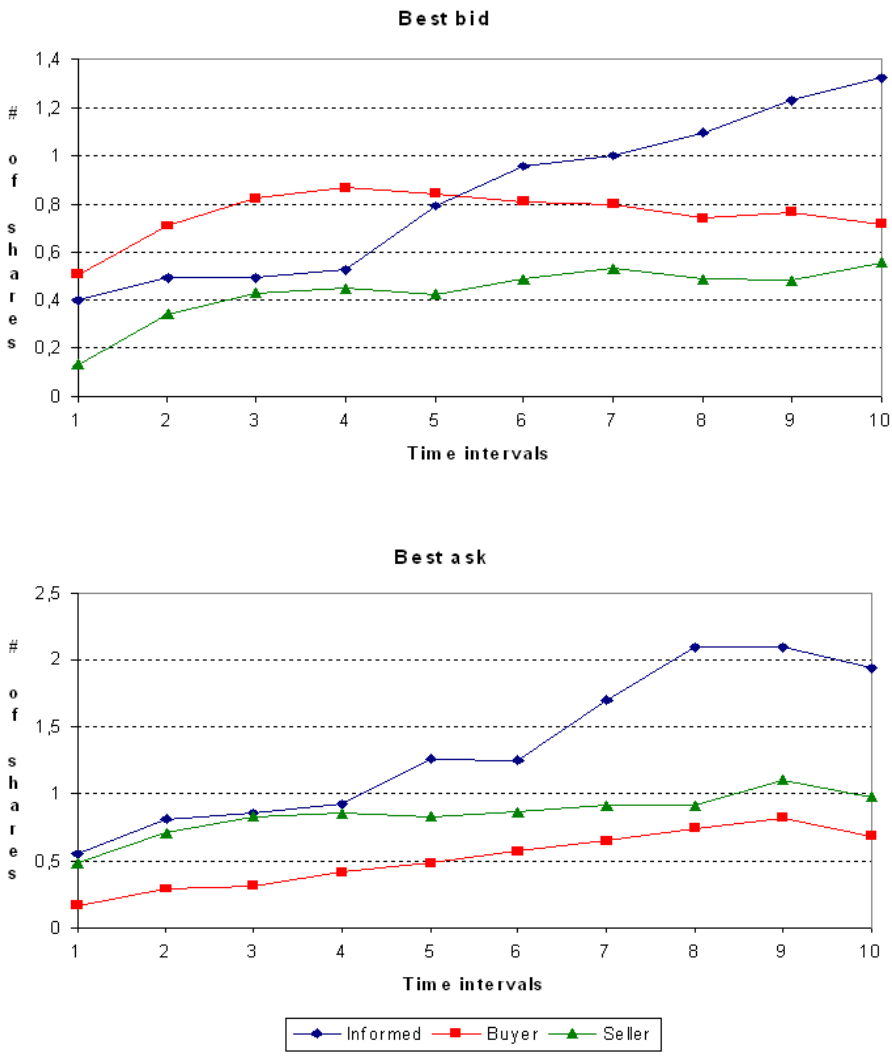

Figure 9: Market quality in both experimental designs

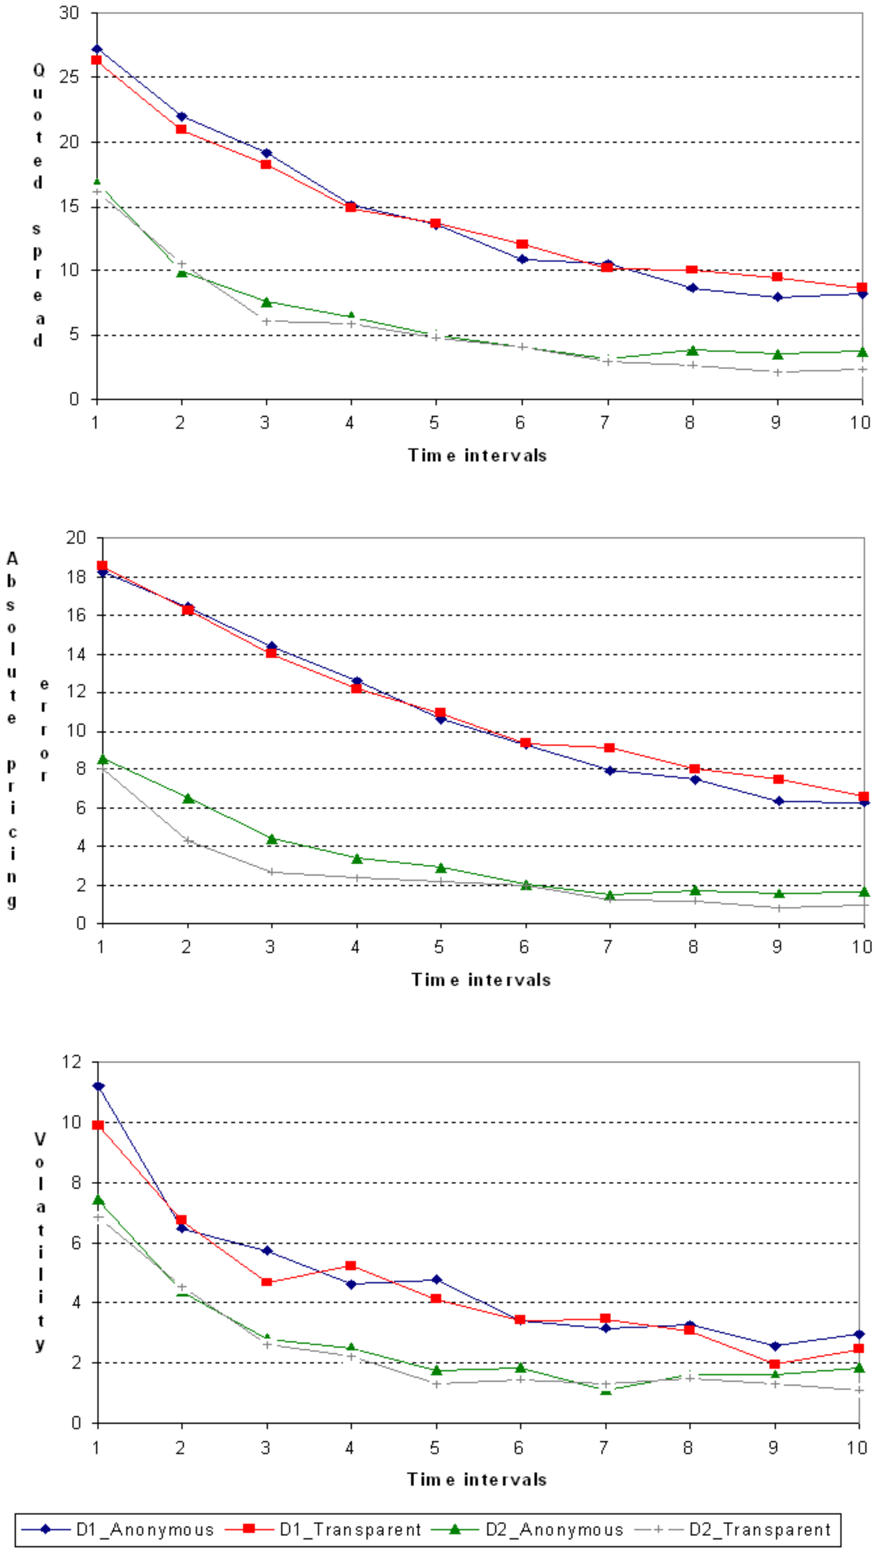

Table 1: Multivariate analysis of the spread

	Constant	Period	Int	Int^2	Value	$Value^2$	Transp	Adj. R^2
Panel A	34.05^{***}	-0.184	-0.314^{***}	0.0011^{***}	-0.159^{*}	0.0018^{**}	0.483	0.15
Panel B	31.92^{***}	-0.189^{**}	-0.312^{***}	0.0011^{***}	-0.207^{***}	0.0024^{***}	0.612	0.43

This table presents the results of the multivariate analysis of the spread.

Panel A presents estimates for the following equation:
$Spread = \alpha + \beta_1 Period + \beta_2 Int + \beta_3 Int^2 + \beta_4 Value + \beta_5 Value^2 + \beta_6 Transp + \epsilon$
where each observation of the dependent variable *Spread* is the value of the spread in a given session, market and interval. *Period* is the number of the experimental market within a session, that can take values from 1 to 18. Int denotes one of the 10 intervals of a market expressed in seconds (15, 30, ..., 150), and Int^2 is the square value of Int. $Value$ ($Value^2$) is the true security value (square of the true value) that is randomly drawn for each market. $Transp$ is a dummy variable that takes the value 1 for markets played in a transparent setting and 0 otherwise.

Panel B presents estimates for a similar specification, where we have added cohort-specific dummy variables. Coefficients on these variables are not reported to conserve space.

For all coefficients in all panels, the marks ***, ** and * represent significance at respectively the 1, 5 and 10% level.

Table 2: Spread by transparency regime and extremity

	Low extremity	Medium extremity	High extremity
Panel A: final state			
Anonymous	7.5	7.8	8.4
Transparent	9.1	9.3	7.5
Panel B: average			
Anonymous	13.5	13.6	15.5
Transparent	11.4	15.1	15.2

To obtain this table, we first assign each market to one of the six cells defined by the transparency (anonymous or transparent) and extremity (low, medium or high) factors. The extremity category is determined by the absolute difference (AD) between the true security value and the prior expected value ($AD \equiv |Value - 50.5|$). The extremity of a market is *low* if $AD < 10$; it is *medium* if $10 \leq AD < 30$, and it is *high* if $AD \geq 30$.
For each cohort, we compute an average value of the spread in each of the six cells, and those values are averaged across cohorts to provide the figures in the table.
In Panel A, we consider the value of the spread at the last interval of a market. In Panel B, we consider the average value of the spread over a market.

Table 3: Multivariate analysis of the absolute pricing error

	Constant	Period	Int	Int^2	Value	Value^2	Transp	Adj. R^2
Panel A	38.57***	−0.412***	−0.176***	0.0005***	−0.702***	0.0063***	1.167	0.29
Panel B	34.74***	−0.413***	−0.175***	0.0005***	−0.725***	0.0066***	1.226	0.38

This table presents the results of the multivariate analysis of the absolute pricing error.

Panel A presents estimates for the following equation:

$APE = \alpha + \beta_1 Period + \beta_2 Int + \beta_3 Int^2 + \beta_4 Value + \beta_5 Value^2 + \beta_6 Transp + \epsilon$

where each observation of the dependent variable APE is the value of the absolute pricing error in a given session, market and interval. $Period$ is the number of the experimental market within a session, that can take values from 1 to 18. Int denotes one of the 10 intervals of a market expressed in seconds (15, 30, ..., 150), and Int^2 is the square value of Int. $Value$ ($Value^2$) is the true security value (square of the true value) that is randomly drawn for each market. $Transp$ is a dummy variable that takes the value 1 for markets played in a transparent setting and 0 otherwise.

Panel B presents estimates for a similar specification, where we have added cohort-specific dummy variables. Coefficients on these variables are not reported to conserve space.

For all coefficients in all panels, the marks ***, ** and * represent significance at respectively the 1, 5 and 10% level.

Table 4: Absolute pricing error by transparency regime and extremity

	Low extremity	Medium extremity	High extremity
Panel A: final state			
Anonymous	4.0	6.3	7.2
Transparent	4.3	5.4	9.3
Panel B: average			
Anonymous	5.6	9.7	15.4
Transparent	5.8	8.7	17.6

To obtain this table, we first assign each market to one of the six cells defined by the transparency (anonymous or transparent) and extremity (low, medium or high) factors. The extremity category is determined by the absolute difference (AD) between the true security value and the prior expected value ($AD \equiv |Value - 50.5|$). The extremity of a market is *low* if $AD < 10$; it is *medium* if $10 \leq AD < 30$, and it is *high* if $AD \geq 30$.

For each cohort, we compute an average value of the absolute pricing error in each of the six cells, and those values are averaged across cohorts to provide the figures in the table.

In Panel A, we consider the value of the absolute pricing error at the last interval of a market. In Panel B, we consider the average value of the absolute pricing error over a market.

Table 5: Traders' profits by transparency and extremity

	Low extremity	Medium extremity	High extremity
Panel A: informed			
Anonymous	16.24	62.44	111.01
Transparent	21.60	59.52	87.82
Panel B: buyers			
Anonymous	-5.68	-30.32	-69.91
Transparent	-13.74	-33.20	-72.28
Panel C: sellers			
Anonymous	-10.56	-32.12	-41.10
Transparent	-7.86	-26.32	-15.54

To obtain this table, we first assign each market to one of the six cells defined by the transparency (anonymous or transparent) and extremity (low, medium or high) factors. The extremity category is determined by the absolute difference (AD) between the true security value and the prior expected value ($AD \equiv |Value - 50.5|$). The extremity of a market is *low* if $AD < 10$; it is *medium* if $10 \leq AD < 30$, and it is *high* if $AD \geq 30$.
For each cohort, we compute an average value of traders' profits in each of the six cells, and those values are averaged across cohorts to provide the figures in the table.
Panel A provides results for informed traders, Panel B for buyers and Panel C for sellers.

Table 6: Buyers' and sellers' profits for extreme security values

	Value ≤ 20	Value > 80
Panel A: buyers		
Anonymous	-103.46	-11.19
Transparent	-117.19	-8.88
Panel B: sellers		
Anonymous	-16.54	-84.08
Transparent	12.67	-55.35

To obtain this table, we only consider markets where the randomly drawn security value is either extremely low (≤ 20) or extremely high (> 80). We assign each of those markets to one of the four cells defined by the transparency (anonymous or transparent) and value (extremely low or extremely high) factors.
For each cohort, we compute an average value of traders' profits in each of the four cells, and those values are averaged across cohorts to provide the figures in the table.
Panel A provides results for buyers and Panel B for sellers.

Table 7: Comparison between designs – multivariate analysis

	Constant	Period	Int	Int^2	Value	Value^2	Transp	Design2	Adj. R^2
Quoted spread	32.15^{***}	-0.236^{***}	-0.300^{***}	0.0011^{***}	-0.110^{*}	0.0015^{**}	0.245	-8.56^{***}	0.21
Absolute error	32.41^{***}	-0.345^{***}	-0.163^{***}	0.0005^{***}	-0.497^{***}	0.0046^{***}	0.696	-8.70^{***}	0.32
Volatility	15.33^{***}	-0.197^{***}	-0.141^{***}	0.0006^{***}	-0.046	0.0005^{*}	-0.116	-3.80^{***}	0.20

This table presents the results of the multivariate analysis for several market quality variables in both experimental designs.

The specification that is estimated has the following form:

$Dep = \alpha + \beta_1 Period + \beta_2 Int + \beta_3 Int^2 + \beta_4 Value + \beta_5 Value^2 + \beta_6 Transp + \beta_7 Design2 + \epsilon$

where each observation of the dependent variable Dep is the value observed in a given session, market and interval. $Period$ is the number of the experimental market within a session, that can take values from 1 to 18. Int denotes one of the 10 intervals of a market expressed in seconds (15, 30, …, 150), and Int^2 is the square value of Int. $Value$ ($Value^2$) is the true security value (square of the true value) that is randomly drawn for each market. $Transp$ is a dummy variable that takes the value 1 for markets played in a transparent setting and 0 otherwise. $Design2$ is a dummy variable that takes the value 1 for markets played in the second experimental design and 0 otherwise.

The three dependent variables under analysis are the quoted spread, the absolute pricing error and volatility. For all coefficients in all panels, the marks ***, ** and * represent significance at respectively the 1, 5 and 10% level.