Over-the-Counter Markets vs Double Auctions: A Comparative Experimental Study*

Giuseppe Attanasi[†] Samuele Centorrino[‡] Ivan Moscati[§]

November 17, 2014

Abstract

In a large-scale classroom experiment we study the functioning of an electronic over-the-counter (OTC) market mechanism, and compare its performance with that of a standard electronic double-auction (DA) mechanism. These two mechanisms differ for an important informational feature: in a DA market traders post their bids and asks publicly, while in an OTC market each agent looks for the best counterpart through bilateral and private bids and asks. Although in many actual markets negotiations occur via computer on a private, bilateral basis, electronic OTC markets have received little attention in the experimental literature. We find that the lack of public information that characterizes our OTC mechanism with respect to a DA mechanism induces a loss of almost eight efficiency points. We also show that this efficiency loss is due to the fact that, in the OTC mechanism, closing prices converge to a price below the competitive price and the traded quantity is lower than the competitive quantity. Finally, we investigate the efficiency and convergence properties of the OTC mechanism when demand or supply shocks modify the competitive equilibrium. Among other things, we show that supply shocks increasing the competitive quantity improve the OTC's efficiency.

JEL classification: C92, D41, D47, D83.

Keywords: Market experiments; Over-the-counter market; Double auction; Private vs. public information; Efficiency.

*This paper has been circulated previously under the title “Double Auction Equilibrium and Efficiency in an Experimental Search Market”. The authors would like to thank Olivier Armantier, Antonio Bisignano, Fortuna Casoria, Luigi Luini, Valentina Rotondi, and participants in the 2011 Labsi Conference, Capua, Italy, the 2012 Annual Xiamen University International Workshop on Experimental Economics and Finance, Xiamen, China, the 2013 Storep Conference, Gaeta, Italy, and the SIE 2013 Conference, Bologna, Italy, for useful comments and remarks. G. Attanasi gratefully acknowledges financial support by the Foundation for an Industrial Safety Culture (Fondation pour une Culture de Sécurité Industrielle, FONCSI) and by the European Research Council (ERC Starting Grant DU 283953). Ivan Moscati thanks Fondazione Cariplo for financial support. The usual disclaimer applies.

[†]BETA, University of Strasbourg, 61 Avenue de la Forêt-Noire, 67085 Cedex, Strasbourg, France. Email: attanasi@unistra.fr

[‡]Economics Department, Stony Brook University, Social and Behavioral Sciences Building, Stony Brook, NY 11774. Email: samuele.centorrino@stonybrook.edu

[§]University of Insubria, Department of Economics, Via Monte Generoso 71, 21100 Varese, Italy. Email: ivan.moscati@uninsubria.it

1 Introduction

Experimental economists have studied the functioning and equilibrium properties of different market institutions for more than half a century. In particular, they have devoted a great deal of attention to the *double-auction* (henceforth *DA*) market introduced by Vernon Smith (1962). In a *DA* market, buyers and sellers trade a single homogeneous good. Buyers can submit public bids for the good and are free to accept asks from sellers, while sellers can submit public asks and are free to accept bids from buyers. When a buyer accepts an ask, or a seller accepts a bid, a public transaction takes place at the accepted price, and both the bid and ask are removed from the market.

Because different units of the commodity are (typically) traded at different prices, and traders are price makers, *DA* markets are not competitive markets. However, Smith (1962) found that transaction prices and exchanged quantity quickly converge to the competitive price and quantity, and the efficiency reached by *DA* markets closely approximates that reached by competitive markets. Moreover, as shown also by subsequent experimental research, the convergence and efficiency properties of *DA* markets are robust to modifications in the number of buyers and sellers, in their redemption values for the good, as well as in the number of units of the good they can buy or sell (for a review of the experimental research on *DA*, see Friedman and Rust, 1993).

Since the late 1970s, Smith and other experimental economists have explored the question of whether the convergence and efficiency properties displayed by *DA* markets are shared by other mechanisms. In particular, this literature has focused on other auction mechanisms such as the *posted offer auction*, in which sellers submit asks and buyers can purchase at the posted prices, or the *clearing house auction*, where buyers and sellers submit their bids and asks to a clearing house, and all units are sold at the price that clears the market. The basic outcome of this research is that other auction mechanisms also tend to converge to the competitive outcome, but they are generally less efficient than *DA*, and their convergence and efficiency properties are less robust than those of *DA* to minor modifications in the experimental design (for a review of this literature, see Holt, 1996, Cason and Friedman, 2008).

In the present paper we explore the convergence and efficiency properties of a trading mechanism that is not an auction, but is similar to the *DA* mechanism under many other respects. Ours is an *over-the-counter* (henceforth *OTC*) decentralized mechanism in which, as in the *DA* setting, buyers and sellers are price makers and trade via computer a homogeneous good, different units of the good are typically traded at different prices, and no clearing house exists in the market. However, while in *DA* markets buyers and sellers post their bid quotes and ask quotes publicly, this public-information element is absent from our *OTC* market. In our *OTC* market, each agent looks for the best counterpart through *private* bids and asks submitted by computer, i.e. by making/receiving one electronic quote at a time to/from a single counterpart. If the quote is accepted, the transaction is carried out, and the closing price appears on the screens of all traders in the market, and is thus made public, as in *DA* markets. This feature of the market may be labeled as *post-trade price transparency*. However, and differently from the *DA* setting, the history of bids and asks that led to that transaction remains private information. Also if the quote is withdrawn, and again in contrast to the *DA* setting, the quote and the fact that it was withdrawn remain private information between the two counterparts.

If regarded in historical perspective, our *OTC* mechanism is very similar to the pit market

designed by Edward Chamberlin (1948) in a seminal contribution to the experimental literature on market institutions. Chamberlin implemented a classroom market where students could roam freely around the room and engage in bilateral and private bargaining (for a review of classroom experiments on pit markets, see Holt, 1996). However, there are three main differences between Chamberlin’s setting and ours. First, in Chamberlin’s classroom the bargaining between a buyer and a seller was conducted orally, and buyers and sellers physically close to the bargainers could learn their bids and asks. In our setting, these informational spillovers are ruled out by the fact that traders interact via computer, and the screen and keyboard of each subject’s computer cannot be seen by other subjects. Second, Chamberlin did not always make public the price of closed transactions, while we always implement post-trade price transparency. Third, Chamberlin let experimental subjects trade for one single market period while we follow Smith (1962) and subsequent standard practice in experimental economics, and allow experimental subjects to trade for several periods so that they can gain experience about how the trading mechanism works.¹

It is important to explore this *OTC* mechanism for it has significant economic applications. In a large number of actual markets, negotiations and transactions occur on a bilateral basis rather than, as happens in auction markets, through publicly posted bids and asks. Moreover, these negotiations and transactions occur via computer rather than, as happens in pit markets, orally. Many types of government and corporate bonds, real estate, currencies, and bulk commodities are typically traded electronically over the counter. In a number of these *OTC* markets, such as those for U.S. corporate and municipal bonds, financial regulators have mandated post-trade price transparency, often implemented through a program called Trade Reporting and Compliance Engine (TRACE) (for a thorough discussion of *OTC* markets, see, e.g., Duffie et al., 2005, Duffie, 2012, Ang et al., 2013).

Despite the economic relevance of *OTC* markets, we are not aware of any other experimental studies that have explored systematically the convergence and efficiency properties of an *OTC* mechanism and compared them to those of the standard *DA* mechanism. Our paper fills this lacuna.

More specifically, in order to investigate whether, in an *OTC* market transaction prices and exchanged quantity converge to their competitive levels, and how the efficiency reached by an *OTC* mechanism compares to that reached by the *DA* mechanism, we ran a series of classroom experiments. The experiments involved more than 3300 undergraduate students of almost the same age over a period of six years, namely from 2009 to 2014. As it is common in most classroom experiments, especially those involving markets (see Cason and Friedman, 2008), we did not use monetary rewards to incentivize the students.

¹Recently, a trading mechanism in the spirit of Chamberlin has been investigated by List (2002, 2004) in field experiments involving a sports card market and a collector pin market. Like in Chamberlin’s setting and differently from ours, in List’s the buyer-seller bargaining is conducted orally rather than via computer. Like us and differently from Chamberlin, however, List allows subjects to trade for multiple periods (four), rather than for a single period. One key feature of List’s experimental design is that in it subjects choose endogenously their role as buyers or sellers, while we follow Chamberlin and Smith in assigning subjects to one of the two roles exogenously and randomly. More generally, the focus of List’s experiments is to examine how the experience of buyers and sellers influence the outcomes of a market *à la* Chamberlin. Our goal, in contrast, is to compare the market outcomes of two trading mechanisms – *OTC à la* Chamberlin and *DA à la* Smith – under the assumption that traders have similar market experience.

Our main research hypothesis was that the information disadvantage of the *OTC* mechanism, where only closing prices are made public, over the *DA* mechanism, where the entire history of bids and asks is public information, makes *OTC* markets less efficient than *DA* markets. Our experimental findings validate this research hypothesis: the *OTC* market is less efficient than the *DA* market. We take as an index for efficiency the ratio between the surplus actually realized from trade and the equilibrium surplus. We find that in *DA* markets the average efficiency index is about 93 over 100, while in *OTC* markets the efficiency index is about 85. Thus the information gap between the *OTC* and the *DA* settings determines a loss of efficiency of almost 8 efficiency points. We show that this result is robust to both subjects' learning and to a reduction in the time length of trading periods.

To better understand how the lack of information about the history of bids and asks affects negatively the efficiency of the *OTC* mechanism, we study the pattern of closing prices and traded quantity in both the *OTC* and the *DA* setting. We find that, because of its informational features, in the *OTC* mechanism closing prices converge to a price that is below the competitive price. This, in turn, implies that the traded quantity is lower than the competitive quantity, which is the main source of the *OTC*'s inefficiency.

We bring the analysis further: we decompose the loss of efficiency associated with both the *OTC* and *DA* mechanisms into two main components – intra-marginal inefficiency and extra-marginal inefficiency – and show that, while the inefficiency associated with the *DA* mechanism is almost completely of the extra-marginal type, the inefficiency of the *OTC* mechanism is an even mixture of both types.

Finally, to deepen our comprehension of the *OTC* mechanism, we introduce shocks into the picture and study how efficiency in the *OTC* and the *DA* mechanisms is affected by different types of shocks, that is, by shifts in either the demand curve or the supply curve that modify the competitive equilibrium. We find that, in the short run, none of these shocks substantially affect the efficiency of either the *OTC* or the *DA* mechanism. Conversely, in the long run, a shock that shifts downwards the supply function is able to significantly increase the efficiency of the *OTC* market only. This is due to a reduction in the difference between the competitive price and the average closing price, this difference being usually positive in the *OTC* market.

2 Experimental Design

Procedures. We ran computerized classroom experiments through the z-Tree software (Fischbacher, 2007). Sessions were held at Bocconi University, Milan, during a first-year introductory course in Microeconomics over six consecutive academic years, from 2009 to 2014. All classroom experiments were held in the month of October (first semester), always in the same computerized room, and run by the same experimenter (G. Attanasi), who is also one of the authors of this paper. About one third of enrolled students per year were involved in the experiments, i.e. 3366 students as a whole. The six cohorts of participants have been homogeneous in many features: age (almost all students being 19 or 20 years old), gender (45% female), nationality (around 80% Italians), and field of study (all were students in Economics). We kept the number of traders

essentially constant (40 subjects) across the 84 experimental sessions we ran.² In 42 (i.e., half) of the sessions the *DA* treatment was implemented (1686 subjects in total), while in the remaining 42 sessions the *OTC* treatment was implemented (1680 subjects in total).

Common Features. Here we describe features of the design common to each of the 84 sessions:

- *Number and length of trading periods.* Each experimental session consists of nine trading periods. The nine periods are partitioned into three phases, with each phase consisting of three periods. The first six periods have equal clock time length, namely 120 seconds per period, hence the first and the second phase last 360 seconds each, that is, 6 minutes per phase. The last three periods (third phase) have a shorter clock time length, namely 60 seconds each, that is 3 minutes in total for the third phase.³
- *Market structure.* The 40 subjects in each session are divided equally into buyers and sellers (20 buyers and 20 sellers for each session). As in Cason and Friedman (1996), subjects are allowed to trade only one unit per period. During each period only one unit of a homogeneous good can be bought/sold by a specific buyer/seller. In particular, every seller owns only one unit of the good. Each buyer (seller) is assigned a valuation (a cost) for the single unit of the good he/she has to buy (sell). The buyer's valuation sets the maximum amount he/she can spend for one unit of the good, while the seller's cost sets the minimum amount he/she has to cash in for his/her unit. As in Smith (1962), valuations and costs are exogenously given. In particular, valuations and costs are distributed so that each buyer (seller) has a different valuation (cost) from those of all other buyers (sellers). By sorting individual valuations from the highest to the lowest, and costs from the lowest to the highest, we obtain a demand and a supply curve, respectively. The competitive-equilibrium price and quantity are determined by the intersection of these two curves (see Figure 1). In particular, to check that the experimental outcomes are independent of the initial conditions, we implemented three different distributions of valuations/costs, leading to equilibrium quantity-price combinations *A*, *B* and *C* in Figure 1.⁴
- *Budget constraints.* Two budget-like constraints are imposed. A feasibility constraint imposes that buyers cannot bid over their own valuation, and sellers cannot ask under their own cost. An intertemporal constraint dictates that wealth cannot be transferred through different periods, hence a buyer cannot use in the next periods the amount not spent in the current period, and a seller cannot sell in the next periods the unit of the good not sold in the current period.

²To be precise, in 81 sessions we had 40 subjects and in 3 sessions we had 42 subjects.

³As will be clearer below, the reason for introducing a shorter third phase is twofold: check whether experimental outcomes are robust to further learning on the part of the subjects, and whether lack of time influences experimental outcomes in a different way in different treatments.

⁴Equilibrium combination *A*, $(q^*, p^*) = (17, 64)$, is obtained from a distribution of valuations v and costs c with $\max v = 98$, $\min c = 32$ and a 2-integer distance between two subsequent valuations or costs, i.e. $v \in \{98, 96, \dots, 62, 60\}$ and $c \in \{32, 34, \dots, 68, 70\}$. Similarly, equilibrium combination *B*, $(q^*, p^*) = (14, 70)$, is obtained with $\max v = 98$, $\min c = 44$ and a 2-integer distance between two subsequent valuations or costs; equilibrium combination *C*, $(q^*, p^*) = (14, 58)$, is obtained with $\max v = 86$, $\min c = 44$ and same 2-integer distance.

Figure 1: Competitive-equilibrium quantity and price for different distributions of valuations/costs, and after different shocks.

- *Information.* At the beginning of each phase, subjects are informed that the phase is constituted by three periods. In each period, subjects do not know the distributions of valuations and costs in the market.⁵ At the beginning of each phase, each subject is given three pieces of information: his/her role (either buyer or seller), his/her redemption value (either valuation or cost) for the single unit of the good, and his/her ID. These are private information and always appear on the subject's screen. It is common knowledge that while subjects' roles and redemption values are kept constant in the three periods of the phase, their IDs are reshuffled at every period. This prevents subjects from identifying trading counterparts in a given period on the basis of IDs learned in previous periods. At the end of the each phase, roles are kept constant while redemption values are reshuffled. Therefore, in the second (third) phase a given subject might have a different redemption value than in the first (first and second) phase (phases). To prevent repeated-game effects, subjects are informed about the existence of a second (third) phase only at the end of the first (second) phase.
- *Incentives.* At the end of each period, each subject sees on the screen his/her payoff as the difference between valuation and closing price – if he/she is a buyer – or between closing price and cost – if he/she is a seller. If a subject does not trade his/her commodity unit within the period, his/her payoff is equal to zero. As it is common in many classroom experiments, and because the exceptionally large number of students involved in our setting would make paying them too expensive, we do not use monetary incentives. However, we

⁵Indeed, subjects are only implicitly told that the support of the valuations and costs is constituted by integer numbers between 0 and 99. This is because the experimental software only allows players to enter numbers with one or two digits in the bid/ask box.

give students a non-monetary incentive to play fairly: at the end of each of the two phases, subjects are ranked according to their corrected total profit in that phase.⁶ We then asked the four subjects having earned the highest total profit in that phase to stand up and praised them publicly for their performance.⁷

Main Treatments. The main treatment variable is the *trading mechanism* used to allow buyers and sellers to interact within a trading period. We have two trading mechanisms, *DA* and *OTC*, with 42 experimental sessions per treatment. In both treatments, once a subject's quote is accepted by a counterpart, the closing price appears on the screens of all subjects, not only of the two traders. In this way prices of closed transactions are made public in chronological order. However, two main features distinguish the *DA* from the *OTC* treatment:

- *Public vs. private quotes.* We define a buyer's (seller's) bid (ask) as "public" if it can be addressed to all sellers (buyers) in the market and it is disclosed to all buyers and sellers (also to those no more in the market). Conversely, we define a quote as "private" if it can be addressed to only one counterpart in the market and only this subject can observe it. In particular:
 - In the *DA* treatment, buyers and sellers post their bids and asks publicly, so that the bid-ask history of the market is public information. Hence, every buyer (seller) is always informed about the best bid (ask) on the market.
 - In the *OTC* treatment, each subject looks for the best counterpart through private bids and asks. More precisely, a subject can send only one quote at a time to a single counterpart, by indicating the amount of the quote and the counterpart's ID. If the counterpart does not reply, the quote may be withdrawn and a new quote can be made that differs either in terms of the amount, the counterpart's ID or both.⁸
- *Bid/ask improvement rule.* The bid/ask improvement rule imposes that, in order to make a valid quote, a subject has to improve on the existing situation. A buyer has to submit a bid higher than the current highest bid (ascending auction), and a seller has to submit an ask lower than the current lowest ask (descending auction). When a buyer and a seller reach an agreement, they exit the market, the standing bids and asks are removed, and new bids and asks can be submitted.
 - In the *DA* treatment, the *bid/ask improvement rule* holds. Every subject is informed about the highest bid and the lowest ask existing in the market, and a transaction

⁶Profits are in fact corrected: since redemption values are assigned randomly, subjects who are less lucky would be penalized. Therefore, we implement a correction factor that, for buyers, is proportional to the distance between their valuation and the highest valuation in the market and, for sellers, is proportional to the distance between their cost and the lowest cost in the market. Before the beginning of the experiment, subjects are informed about the way profits will be corrected.

⁷On the methodology of classroom experiments and the issue of whether monetary incentives are really necessary to motivate experimental subjects, see Holt (1999), Guala (2005), and Bardsley et al. (2009).

⁸Notice that it is possible that a subject receives more than one offer at a time (because his/her ID has been indicated by more than one counterpart during the same time interval). In this case, offers are automatically ranked, so that the best possible deal always appears on the top of the subject's screen.

between a buyer and a seller is realized when either the former accepts the standing lowest ask of the sellers' descending auction or the latter accepts the standing highest bid of the buyers' ascending auction. The closing price is the accepted quote.

- In the *OTC* treatment, the *bid/ask improvement rule* does not hold: subjects do not observe the best bid and ask present in the market, and, whenever they withdraw a quote, they can replace their previous bid (ask) with a lower (higher) one in the new quote.

Additional Treatments. We define an *exogenous shock* as a modification in the buyers' valuations (i.e., a shift in the market demand curve) or a modification in the sellers' costs (i.e., a shift in the market supply curve) that leads to a change in the competitive quantity q^* and the competitive price p^* . Further treatment variables are: the possibility that an exogenous shock is introduced in the second phase; whether this shock concerns the buyers' valuations or the sellers' costs; whether valuations (or costs) are increased or decreased.

- *No shock vs. shock.* In 36 over 84 experimental sessions no shock is applied in the second phase. In all phases of the experiment, the distribution of valuations/costs is always the same, leading to either *A*, or *B* or *C* in Figure 1 (6 experimental sessions per treatment for each quantity-price combination).⁹ In the remaining 48 experimental sessions a shock is applied to either demand or supply for both the *DA* treatment (24 sessions) and the *OTC* treatment (24 sessions). In all these sessions, a shock occurs at the beginning of the second phase (period 4), and is maintained during the whole second phase (periods 4–6) and during the third phase (periods 7–9). The third phase is mainly run in order to check whether further learning on the part of the subjects may have different effects on the experimental outcomes when a specific shock is applied.
- *Types of shocks.* We implement four different types of shocks (6 experimental sessions for each type of shock, per treatment). Each type of shock is characterized by two features: whether the variation concerns the support of the valuations or the support of costs; whether all redemption values in a support are increased or decreased by the same amount.¹⁰ Shocks produce shifts in either the demand or the supply curve and thus lead to a change in the predicted competitive quantity q^* and competitive price p^* . By defining as *negative (positive)* a shock that leads to an decrease (increase) of the competitive quantity q^* , the four shocks can be classified as:¹¹

1. a negative (downward) shift of demand, indicated as D^- and leading to a decrease of both q^* and p^* (in Figure 1, from *A* to *C*);

⁹In particular, we have 12 no-shock sessions with *A* as predicted equilibrium combination in all trading periods (6 under *DA* and 6 under *OTC*); 12 no-shock sessions with *B* (6 under *DA* and 6 under *OTC*); 12 no-shock sessions with *C* (6 under *DA* and 6 under *OTC*).

¹⁰This amount is 12 integers for each shock; i.e., it is the distance between the two demand functions in Figure 1, that is equal to the distance between the two supply functions in the same figure.

¹¹Given that two of the four types of shock have *A* as pre-shock predicted combination, of the 24 sessions with shocks for each main treatment (*DA* and *OTC*), 12 have predicted combination *A*, 6 have predicted combination *B*, and 6 have predicted combination *C*.

2. a positive (upward) shift of demand, indicated as D^+ and leading to an increase of both q^* and p^* (in Figure 1, from combination C to combination A);
 3. a negative (upward) shift of supply, symbolized by S^- and determining a decrease of q^* and an increase of p^* (in Figure 1, from A to B);
 4. a positive (downward) shift of supply, symbolized by S^+ and determining an increase of q^* and a decrease of p^* (in Figure 1, from B to A).
- *Information.* Subjects are given no information about the fact that a shock has been applied in the second phase and maintained in the third phase, nor about the type of shock.

Table 1 summarizes our experimental design by indicating the absolute number of sessions that we ran for each market mechanism (main treatments: DA and OTC), without shock and for each of the four types of shock (D^- , D^+ , S^- , S^+). As Table 1 shows, we ended up having 10 treatments, each one characterized by a different mechanism-shock combination.

	No Shock	D^-	D^+	S^-	S^+
DA	18	6	6	6	6
OTC	18	6	6	6	6

Table 1: Number of sessions per treatment.

3 Results

The basic measure on which we rely to compare the performance of DA and OTC mechanisms is the *efficiency index* defined by Smith (1962) and used, among others, by Gode and Sunder (1993). It is the ratio between the the total uniperiodal profit actually earned by all traders in a market (i.e., the sum of consumer and producer realized market surplus), and the maximum total uniperiodal profit that could have been earned by all traders in the market (i.e., the sum of consumer and producer equilibrium surplus). The efficiency index goes from 0 (minimal efficiency) to 100 (full efficiency), with full efficiency being reached if all subjects trade at the equilibrium price.

The efficiency index, however, says nothing about the causes of the inefficiency. To better understand these causes, we also compare the *quantity actually traded* in these markets with the predicted competitive quantity, and study the *pattern of closing prices* under the DA and OTC mechanisms. Moreover, following Rust et al. (1993), we decompose inefficiency into two main components: *intra-marginal inefficiency* and *extra-marginal inefficiency*. There is intra-marginal inefficiency (henceforth *IM-inefficiency*) when two intra-marginal traders – i.e., a buyer with a valuation higher than the competitive price and a seller with a cost lower than the competitive price – do not exchange. Extra-marginal inefficiency (*EM-inefficiency*), by contrast, occurs when an extra-marginal trader – i.e., a buyer with a valuation lower than the competitive price or a seller whose cost is higher than the competitive price – exchanges with an intra-marginal trader.

First we compare the relative performance of *DA* and *OTC* mechanisms in the 36 sessions (18 per mechanism) without shocks. Then we perform the same analysis for the 48 sessions (24 per mechanism) characterized by an exogenous shock in period 4, that is, at the beginning of the second phase. In order to facilitate comparison between the treatments without shocks (section 3.1) and those with shocks (section 3.2), we report results for the entire experimental session (periods 1–9) and for the first phase (periods 1–3), the second phase (periods 4–6), and the third phase (periods 7–9) separately.¹²

3.1 DA and OTC without shocks

Table 2 reports the *efficiency index* for the *DA* and *OTC* markets respectively. Recall that the informational advantage of the *DA* over the *OTC* mechanism is due to the fact that in the former the entire history of bids and asks is public information, while in the latter only the closing prices are made public. As expected, this makes the *DA* market more efficient, on average, than the *OTC* market. More precisely, over periods 1–9 we observe an average efficiency index of 93.1 in the *DA* market: this result is in line with past findings of Gode and Sunder (1993), and Cason and Friedman (1996). We also find that the average efficiency index is greater in the *DA* than in the *OTC* market (85.3), with this difference being significant at 1% level (see results of Mann-Whitney test in the first row, first column of Table 7 in the Appendix, rejecting the null hypothesis of equal population medians with $P\text{-value} = 0.000$). This significant loss of efficiency in the *OTC* market is quantified in 7.8 efficiency points on average across all trading periods.

	DA	OTC
Period 1	87.8	68.3
Period 2	94.7	84.8
Period 3	93.3	84.6
Phase 1	91.9	79.2
Period 4	95.2	92.1
Period 5	94.2	92.7
Period 6	93.9	89.6
Phase 2	94.4	91.5
Period 7	93.9	91.6
Period 8	93.0	83.6
Period 9	92.4	80.1
Phase 3	93.1	85.1
Total	93.1	85.3

Table 2: Efficiency index in the no-shock treatment.

As it has been already shown in other experimental studies (see Cason and Friedman, 1996, and references therein), experience in market experiments, i.e. learning, increases the efficiency of a market institution. In our study this is true, in particular, for the *OTC* mechanism. Over the first phase of the experiment (periods 1–3) its average efficiency index is in fact 79.2, which

¹²All data and experimental instructions are available upon request.

is more than 10 points smaller than the average efficiency index of the *DA* mechanism (91.9) and more than 20 points far from full efficiency. This situation leaves more room for efficiency gains due to traders' learning. In fact, in passing from the first to second phase of the experiment (periods 4–6), the average efficiency of the *OTC* mechanism increases by more than 12 efficiency points (from 79.2 to 91.5) and the gap with the average efficiency of the *DA* mechanism (94.4) reduces to just 2.9 efficiency points. However, the improvement in the efficiency of the *OTC* mechanism disappears when the trading time shortens. In the third phase of the experiment (periods 7–9), when the trading time in each period is only 60 seconds, that is half the time of all previous periods, the efficiency index of the *OTC* mechanism significantly decreases, and passes from 91.5 (second phase) to 85.1 (third phase). This loss of efficiency appears to be due to the lack of time to privately look for the best counterpart in the *OTC* market. This supposition is confirmed by the fact that in the *DA* market, which is less affected by lack of time thanks to the faster public double-auction trading procedure, the average efficiency in the third phase (93.1) is only slightly smaller than the average efficiency in the second phase (94.4). It is important to notice, however, that the average efficiency index for the *OTC* market in the third phase of the experiment (85.1) is significantly greater (6 efficiency points) than in the first phase (79.2). Therefore, it seems that in the *OTC* market learning effects partially compensate the increased difficulty in finding the best counterpart when the time available is halved.

The above findings are summarized as follows:

Result 1 (Efficiency). *The DA market is significantly more efficient than the the OTC market. The efficiency gap can be quantified in 7.8 efficiency points over all trading periods. Learning partially offsets this efficiency gap, although this effect is smaller the smaller the length of the trading period.*

Now we focus on the analysis of the causes of the higher inefficiency of the *OTC* mechanism. This may come from two sources: the actual traded quantity q being different from the competitive quantity q^* ; and/or closing prices converging to a price that is different from the competitive price p^* .

Table 3 reports, for each of the nine trading periods, the fraction of sessions where the *traded quantity* is different than the competitive one. Indeed, if $q < q^*$, then profitable trades between some intra-marginal buyer and some intra-marginal seller have not taken place. If $q > q^*$, then some commodity units that should have been left out of the market have instead been exchanged: either some extra-marginal buyer managed to buy from an intra-marginal seller, or some extra-marginal seller managed to sell his/her unit to an intra-marginal buyer, or both. Notice that exchange between two extra-marginal traders is impossible: extra-marginal buyers have valuations below the competitive price, while extra-marginal sellers have costs above the competitive price. Therefore, the set of possible agreements between these two categories of traders is empty.

Table 3 shows that in the *DA* market, $q < q^*$ in only 14% of all trading periods. This result is consistent with existing experimental evidence about *DA* (see, e.g., Gode and Sunder, 1993, Cason and Friedman, 1996), and appears to be due to the informational features of the *DA* mechanism. Since in these markets the current highest bid and the current lowest ask are public information, it is easy for intra-marginal buyers (who have higher valuations than extra-marginal buyers) and intra-marginal sellers (who have lower costs than extra-marginal sellers) to propose

	DA			OTC		
	$q < q^*$	$q = q^*$	$q > q^*$	$q < q^*$	$q = q^*$	$q > q^*$
Period 1	33.3	33.3	33.4	100.0	0.0	0.0
Period 2	0.0	38.9	61.1	55.6	33.3	11.1
Period 3	11.1	38.9	50.0	44.4	27.8	27.8
Phase 1	14.8	37.0	48.1	66.7	20.4	13.0
Period 4	5.6	22.2	72.2	33.3	27.8	38.9
Period 5	11.1	27.8	61.1	22.2	38.9	38.9
Period 6	11.1	27.8	61.1	50.0	22.2	27.8
Phase 2	9.3	25.9	64.8	35.2	29.6	35.2
Period 7	11.1	22.2	66.7	44.4	27.8	27.8
Period 8	11.1	50.0	38.9	72.2	11.1	16.7
Period 9	27.8	27.8	44.4	72.2	22.2	5.6
Phase 3	16.7	33.3	50.0	63.0	20.4	16.7
Total	13.6	32.1	54.3	54.9	23.5	21.6

Table 3: Traded quantity q vs. competitive-equilibrium quantity q^* in the no-shock treatment (for each mechanism, in percentage over all sessions within the same period).

deals that can be accepted by an intra-marginal counterpart. This aspect of the *DA* mechanism is not significantly affected by learning or the duration of the trading period, as the percentage of periods for which $q < q^*$ remains quite stable over the three phases of the experiment.

In the *OTC* market, $q < q^*$ much more frequently than in *DA* market, namely in 55% of all trading periods. Our explanation for this result is that, since in *OTC* markets negotiations are conducted on a one-to-one basis, intra-marginal traders can easily miss the possibility of closing a profitable transaction before the end of the trading period, due the limited trading time. To support this explanation, notice that, when traders become more experienced, the percentage of periods for which $q < q^*$ strongly decreases: it is 67% on average in the first phase, and only 35% on average in the second phase. This corresponds to a notable increase in the efficiency of the *OTC* mechanism (+12.3 efficiency points), as reported in Table 2 above. Correspondingly, the decrease in efficiency from the second to the third phase in the *OTC* market (−6.4 efficiency points) is associated with a significant increase (from 35% to 63% on average) in the percentage of periods for which $q < q^*$. In the *OTC* mechanism, therefore, the positive effects of learning on the quantity traded disappear as the length of the trading period decreases.

These findings can be summarized as follows:

Result 2 (Traded quantity). *The DA trading mechanism only rarely delivers a traded quantity q lower than the competitive quantity q^* . This is independent of traders’ learning. Conversely, the OTC trading mechanism very often delivers $q < q^*$. Learning about the trading mechanism significantly reduces the number of trading periods where $q < q^*$, although this effect is smaller the smaller the length of the trading period.*

The result that in the *OTC* market $q < q^*$ can be related to the *pattern of closing prices*. If the market converges to an average closing price below (above) the competitive price p^* , intra-marginal sellers (intra-marginal buyers) – who in a competitive market would have sold (bought)

a unit of the good – are left out of the market. This, in turn, reduces the quantity exchanged and generates inefficiencies. We therefore take a closer look at the pattern of closing prices and at their (possible) convergence to the competitive equilibrium.

In the left panel of Figure 2, we draw the pattern of closing prices in all trading periods of a single *DA* experimental session, while in the right panel we do the same for a single *OTC* session (as we will see in a moment, what happens in these two specific sessions is representative of what occurs in all other sessions) For each of the 9 periods of both sessions, the predicted equilibrium combination is A in Figure 1, i.e. $(q^*, p^*) = (17, 64)$. In both panels of Figure 2, the units traded are plotted in abscissas according to the chronological order in which they have been traded in each period (the first unit traded in period 1 comes first on the abscissas line, the second unit traded in period 1 comes second, etc.), with units traded in period 2 plotted after those traded in period 1, and so on until period 9. The prices corresponding to each traded unit are plotted in ordinates; the dashed line corresponds to the predicted competitive price, while the continuous line expresses the average closing price over all 9 periods.

Figure 2 shows that price convergence to the average closing price occurs in both markets. However, in the *DA* treatment the average closing price almost coincides with the competitive price (the continuous and dashed lines are superposed). In contrast, in the *OTC* treatment the average closing price is clearly below the competitive price. The latter result might be explained by the fact that, despite the experimentally-imposed symmetry between the role of buyer and the role of seller, sellers feel much more pressure than buyers in finding a trading counterpart. They want to get rid of the unit of the good they own, and thus are willing to sell it cheaply (see Feldhütter, 2012).

As a matter of fact, the average share of the total surplus allocated to buyers in the *DA* market is 54.8%; while it is equal to 60% in the *OTC*. This difference is significant at 1% level using a Wilcoxon rank sum test.

This pattern of closing prices does not characterize only the specific *DA* and *OTC* experimental sessions represented in Figure 2, but holds for all the *DA* and *OTC* experimental sessions without shocks. This is shown in Figure 3, where we report the histogram of the relative deviations from the competitive price for periods 4–6 of all experimental sessions with no shock. Recall that (see Table 2), among the three phases of the experiment, the highest efficiency is found in the second phase (periods 4–6) and that this holds for both trading mechanisms. Given that the maximum (full) efficiency is found by construction when $(p, q) = (p^*, q^*)$, we guess that if convergence to the competitive equilibrium were reached, this would happen in the second phase. Therefore, we plot the histogram of $\Delta p_i^* = (p_i - p^*)/p^*$, where p_i is the closing price of commodity unit i in periods 4–6 and p^* is the competitive price.¹³ In the left panel of Figure 3, we report the histogram of Δp_i^* for the *DA* treatment: Although this distribution is slightly skewed on the left, it is roughly centered around 0. This confirms that, over *all* *DA* sessions with no shock, the closing prices converge to p^* in the second phase of the experiment. In contrast, the histogram for the *OTC* market – Figure 3, right panel – shows that relative deviations of closing prices are almost normally distributed and that their mean is slightly below 0. This confirms

¹³Notice that, for each trading mechanism (*DA* and *OTC*), we do not find significant differences in experimental outcomes across the three different distributions of valuations/costs (i.e., among A , B and C in Figure 1, for the 18 sessions without shocks in each main treatment). Therefore, without loss of generality, given a trading mechanism, we can analyze Δp_i^* by pooling data of all sessions with no shock.

Figure 2: Closing price patterns in a *DA* market (left panel) and an *OTC* market (right panel) in a no-shock treatment. Dashed line: competitive price; continuous line: average closing price.

that, even when we consider *all OTC* experimental sessions, the closing prices converge to a price below the competitive price. This, in turn, implies that some intra-marginal sellers are left out of the market, that the exchanged quantity remains lower than the competitive quantity, and that inefficiencies emerge.

The findings summarized in Figure 1 and in Figure 3 are also confirmed by the dynamic panel data regression of closing prices in Table 8 in the Appendix. In both main treatments *DA* and *OTC*, autocorrelation coefficients of lag prices are all positive and lower than 1, which confirms price convergence under both trading mechanisms. Furthermore, in the joint regression, the *OTC*-treatment dummy is significant and negative, thereby confirming that closing prices are on average lower in the *OTC* than in the *DA* treatment.¹⁴

Therefore, the following result can be stated:

Result 3 (Closing price). *In the DA treatment the average closing price almost coincides with the competitive price, while in the OTC treatment the average closing price is significantly lower than the competitive price. All this also holds in the second phase of the experiment, where closing prices are more likely to converge to the equilibrium price.*

Result 3 is further explored in a regression analysis of efficiency that is based not only on the data for periods 4–6, but on data of all periods (see Table 9 in the Appendix).¹⁵ This regression

¹⁴Since we have a large panel of price processes for each period, we use the System GMM estimator developed by Blundell and Bond (1998). We run a regression for each treatment separately and then a joint regression with both treatments.

¹⁵We use a zero-or-one inflated beta regression model, since the dependent variables are bounded in the interval $[0, 1]$ and can take values at the boundaries with positive probability. For a general approach and an extensive

Figure 3: Histograms of the absolute relative deviation of closing prices from the equilibrium price, periods 4–6 in *DA* (left panel) and *OTC* (right panel), in no-shock treatment.

shows that efficiency in *DA* markets is increasing as the average trading price $\bar{p}$ approaches the equilibrium price, p^* , that is, as the absolute difference $|\bar{p} - p^*|$ shrinks. By contrast, in the *OTC* market a decrease in $|\bar{p} - p^*|$ does not lead to any statistically significant change in efficiency. This finding is consistent with our descriptive analysis about the two markets. The *DA* market exhibits a very consistent behavior, in that $\bar{p}$ is often close to p^* (for about 40% of our trading periods, the absolute difference is lower than 1 unit). Therefore, even a very slight change in the average trading price can entail a change in the efficiency for this market. However, in the *OTC* mechanism this difference between $\bar{p}$ and p^* is often large (for more than 76% of the trading periods $\bar{p}$ is more than 1 unit away from p^* and, in 74% of them, this difference is negative). Therefore, there is not sufficient variation for the absolute difference to have a significant impact on efficiency.

Notice that this is also consistent with the informational features of the two markets. In the *DA* market, the bid-ask improvement rule drives out intra-marginal players from the game as they can hardly beat the quotes made by extra-marginal players. The pricing feature of the market becomes very important in this context, as slight deviations from the equilibrium price (and for a given equilibrium quantity), can change the total surplus by excluding extra-marginal players or including intra-marginal ones. In the *OTC* market, there is no information about the bids and asks of other players. Therefore, the mass of trades becomes a more crucial determinant of market efficiency. A traded quantity closer to the equilibrium quantity generates by itself more market efficiency in the *OTC* mechanism. Although this may also cause the average trading price to be closer to the equilibrium price, the latter effect seems to go through only via the distance

description of beta regressions see Ospina and Ferrari (2012).

between the traded and the equilibrium quantity.¹⁶

The analysis can be brought further. To better understand the efficiency gap between the *DA* mechanism and the *OTC* mechanism, it is useful to decompose the loss of surplus generated by these two market mechanisms into *IM-inefficiency*, which emerges when two intra-marginal traders do not exchange; and *EM-inefficiency*, which occurs when an extra-marginal trader exchanges with an intra-marginal trader.

There is a somewhat tricky relationship between IM-inefficiency, EM-inefficiency, the traded quantity q , and the competitive quantity q^* . First, IM-inefficiency is linked to a decrease in the traded quantity: if two intra-marginal traders do no exchange, then q cannot be greater than q^* . EM-inefficiency, in contrast, could be linked to an increase in the traded quantity, though it can be present also if q is unchanged. To see why, recall that exchange between two extra-marginal traders is impossible. Thus, an extra-marginal trader always trades a commodity unit with an intra-marginal trader, and in so doing he/she displaces some intra-marginal trader. Two things may happen to a displaced intra-marginal trader: he/she may find another extra-marginal trader with whom he/she trades a commodity unit and in this case q increases; otherwise, he/she may be unable to trade a commodity unit, in which case the quantity q traded on the market does not change. To complicate the picture, when $q < q^*$, both IM-inefficiency and EM-inefficiency can be present. For instance, imagine that intra-marginal buyer $B1$ and intra-marginal seller $S1$ are unable to trade: this decreases the traded quantity and generates IM-inefficiency. However, at the same time, intra-marginal buyer $B2$ trades with extra-marginal seller $S3$ who bumps intra-marginal seller $S2$: This creates EM-inefficiency but does not modify the traded quantity, which remains one unit below q^* .

We can summarize the relationships between IM-inefficiency, EM-inefficiency, q and q^* as follows:

- if $q > q^*$, the only source of inefficiency is EM-inefficiency;
- if $q = q^*$ and realized surplus equals equilibrium surplus, the inefficiency equals 0: all intra-marginal traders have traded their unit and all extra-marginal traders are out of the market;
- if $q = q^*$ but realized surplus is lower than equilibrium surplus, then the existing efficiency is certainly due to EM-inefficiency. IM-inefficiency is ruled out because q is not smaller than q^* ;
- if $q < q^*$, we certainly have IM-inefficiency, but we may also have EM-inefficiency.

Based on our data – that include the redemption values of all traders, the competitive quantity and price, the equilibrium surplus, the quantity actually traded, and the closing prices of all traded units – we can decompose the loss of surplus associated with a specific trading mechanism (*DA* or *OTC*) into its IM-inefficiency and EM-inefficiency components.¹⁷ Table 4 reports the

¹⁶This is also confirmed by the fact that the same regression analysis conducted by omitting the difference in quantity gives a negative and significant coefficient for the difference between $\bar{p}$ and p^* in the *OTC* market.

¹⁷It is worth mentioning that, despite the fact that we are discussing and presenting the aggregate effects of intra-marginal and extra-marginal inefficiency, we perform our analysis recording the instances of each type of inefficiency for each trading period.

results of our market inefficiency audit, by indicating – separately for the *DA* mechanism and for the *OTC* mechanism – the percentage of IM-inefficiency and of EM-inefficiency behind the loss of surplus generated by a specific trading mechanism (notice that, for each mechanism, the sum of the two percentages is 1).

	IM-Inefficiency		EM-Inefficiency	
	DA	OTC	DA	OTC
Period 1	16.6	70.9	83.4	29.1
Period 2	0.0	29.3	100.0	70.7
Period 3	5.5	23.3	94.5	76.7
Period 4	0.0	15.4	100.0	84.6
Period 5	3.1	11.0	96.9	89.0
Period 6	6.1	20.1	93.9	79.9
Period 7	6.6	18.6	93.4	81.4
Period 8	6.8	33.2	93.2	66.8
Period 9	11.8	42.7	88.2	57.3
Total	6.3	26.2	93.7	73.8

Table 4: Sources of inefficiency (in percentage over total inefficiency) in the no-shock treatment.

We find that in the *DA* market IM-inefficiency accounts on average for only 6.3% of total inefficiency, while EM-inefficiency accounts for the residual 93.7% percent. This is congruous with the data displayed in Table 3, which show that in *DA* markets the traded quantity q is lower than the competitive quantity q^* only in 14% of all market periods (recall that we have IM-inefficiency, possibly associated with EM-inefficiency, only when $q < q^*$). In the *OTC* setting, by contrast, we have a mixture of IM-inefficiency and EM-inefficiency. However, as in the *DA* setting, EM-inefficiency plays a more important role: on average across all sessions, IM-inefficiency accounts for only 26.2% of the total market inefficiency. Again, this finding is consistent with the data of Table 3 according to which in the *OTC* setting q is lower than q^* in 55% of all market periods. Furthermore, notice that in the *OTC* setting, the weight of IM-inefficiency is significantly greater than 50% only in trading period 1, where several inexperienced intra-marginal traders are unable to find the best counterpart within the allowed 3 minutes. The highest IM-inefficiency is found in period 1 also in the *DA* setting, although here it only accounts for 16.6% of the loss of surplus. Again, in the *DA* mechanism public asks/bids facilitate immediate learning of the trading mechanism and reduce the time needed to find the best counterpart to trade with. The findings summarized in Table 3 and in Table 4 are confirmed by Beta regression analysis of efficiency and IM-inefficiency by treatment (see Table 9 in the Appendix): in both the *DA* and the *OTC* market, an increase in the traded quantity q with respect to the equilibrium quantity q^* increases efficiency and decreases (increases) IM-inefficiency (EM-inefficiency). All these effects are significant at 1% level.

These findings are summarized in the following:

Result 4 (Sources of inefficiency). *The source of inefficiency in the DA market is almost exclusively extra-marginal. An important amount of intra-marginal inefficiency is instead detected in the OTC market, though significantly smaller than the amount of extra-marginal inefficiency.*

3.2 DA and OTC with shocks

In order to grasp better the functioning of the *OTC* mechanism and compare it with the functioning of the *DA* mechanism, we introduce shocks into the picture and study how the efficiency of each mechanism is affected by shocks. Recall that, in our design, shocks are shifts in either the demand or the supply curve that lead to a change in both the competitive quantity q^* and the competitive price p^* . In all 48 experimental sessions with shocks, the shock occurs in period 4 (first period of the second phase) and is maintained until the end of the experiment, i.e. both during the second phase (periods 4–6) and during the third phase (periods 7–9). We implement four types of shock: D^- , which decreases q^* and p^* ; D^+ , which increases q^* and p^* ; S^- , which decreases q^* and increases p^* ; and S^+ , which increases q^* and decreases p^* .

Table 5 reports the efficiency index for *DA* and *OTC* markets in period 4, over periods 4–6 (second phase), and over periods 7–9 (third phase) in the case without shocks (first column) and for each of the four types of shock.¹⁸

	No shock			D^-			D^+			S^-			S^+		
	Periods			Periods			Periods			Periods			Periods		
	4	4–6	7–9	4	4–6	7–9	4	4–6	7–9	4	4–6	7–9	4	4–6	7–9
DA	95.2	94.4	93.1	90.2	93.1	91.7	96.8	96.2	93.9	91.3	93.6	94.8	97.3	97.6	96.0
OTC	92.1	91.5	85.1	90.0	88.4	86.4	95.7	93.4	86.0	90.0	89.5	86.7	97.3	94.8	89.7

Table 5: Efficiency index by main treatment and type of shock.

Consider first the *DA* market. Without shocks, the efficiency index in period 4 is equal to 95.2. When shocks decreasing q^* are implemented (i.e. D^- and S^-), the efficiency index in period 4 is slightly lower: -5 efficiency points for D^- and -3.9 efficiency points for S^- . In contrast, for shocks increasing q^* (i.e. D^+ and S^+), the efficiency index is slightly higher in period 4: $+1.6$ efficiency points for D^+ and $+2.1$ for S^+ . Over periods 4–6, the efficiency loss associated with negative shocks tends to vanish: for the D^- shock the efficiency index (93.1) approaches the one without shocks (94.4). The same happens for the S^- shock, where the efficiency index (93.6) is even closer to the treatment without shocks. Conversely, the efficiency gap between the baseline treatments and those with positive shocks hangs over ($+1.8$ for D^+ and $+3.2$ for S^+). Table 5 shows similar qualitative results over periods 7–9, coupled with a slight decrease of efficiency – independent of the presence of a shock and of the type of shock – due to trading periods of reduced time length. Notice that, given the sign of the shock, no significant asymmetry is detected between shocks concerning the demand function and shocks concerning the supply function in any of periods 4–9.

The main finding is in accord with those presented in other studies of *DA* markets with shocks (see, e.g. Davis et al., 1993): The temporary efficiency loss provoked in *DA* markets by shocks D^- and S^- may be explained by the fact that both shocks increase the fraction of extra-marginal traders in the market (D^- increases the fraction of extra-marginal sellers while S^- increases the fraction of extra-marginal buyers). This, in turn, increases the probability that extra-marginal traders manage to exchange with some intra-marginal trader and therefore raises

¹⁸Although efficiency in periods 1–3 is an important benchmark to compare efficiency *within* a given treatment, we focus here only on efficiency comparison *between* treatments.

the EM-inefficiency of the *DA* mechanism. However, due to the trading-enhancing features (e.g., public asks/bids) of the *DA* mechanism, negative shocks are absorbed within few periods.

Compared to the *DA* market, in the *OTC* market we observe in period 4 a similarly small efficiency loss due to negative shocks (with respect to 92.1 in the baseline, -2.1 both for D^- and S^-) and a stronger efficiency gain due to positive shocks ($+3.6$ for D^+ and $+5.2$ for S^+). Under S^+ , in period 4 the efficiency of the *OTC* mechanism (97.3) reaches exactly the same level as in the *DA* mechanisms. Thus, we can observe that soon after a shock of any type takes place (period 4), the efficiency of the *DA* and *OTC* mechanisms is very close. Mann-Whitney test in Table 7 in the Appendix does not reject the null hypothesis of equal population medians between efficiency in *DA* and efficiency in *OTC* in period 4 (second column) for each of the four types of shock (D^- , D^+ , S^- , S^+).

Differently from the *DA* market, in the *OTC* market the efficiency gap between the baseline and the negative shocks hangs over during periods 4–6 (-3.1 for D^- and -2 for S^-). As in the *DA* treatment, also the efficiency gap between the baseline and the positive shocks hangs over during periods 4–6 ($+1.9$ for D^+ and $+3.3$ for S^+). Notice that, if the positive shock concerns the supply function, then its positive effect on efficiency hangs over in periods 7–9 too (S^+ vs. no shock: $+4.6$ in periods 7–9). For all other shocks, reduced time length of trading periods 7–9 leads to a sharply decrease in the efficiency index, as it happens in the treatment without shock, with no significant difference in the efficiency index between shock vs. no shock treatments.

The following result can therefore be stated:

Result 5.a (Efficiency after shocks in *DA* and in *OTC*). *Shocks that reduce (increase) the competitive quantity slightly decrease (increase) the efficiency of both the *DA* and the *OTC* market. In the long run, the efficiency gap with respect to the no-shock treatment presents the following trends: (i) it vanishes for both types (demand and supply) of negative shock in both markets; (ii) it hangs over for both types of positive shock in the *DA* market; (iii) in the *OTC* market, it hangs over and increases only for a positive shock S^+ in the supply function.*

Coming back to the comparison between efficiency in the *DA* market and in the *OTC* market, we have shown above that the introduction of a shock changes the main finding of Result 1 in the short run: after a shock, we observe a similar efficiency in the two markets in period 4. However, the main finding of Result 1 holds in the long run: Mann-Whitney test in Table 7 in the Appendix rejects the null hypothesis of equal population medians between efficiency in *DA* and efficiency in *OTC* in periods 4–6 and in periods 7–9 for all types of shock (D^- , D^+ , S^- , S^+). This is summarized in the following:

Result 5.b (Efficiency gap after shocks). *The introduction of a shock offsets the efficiency gap between the *DA* and the *OTC* market in the period where the shock is applied: This is independent of the type of shock. However, the gap emerges again after few periods and increases as the length of the trading period shrinks.*

The specific effect of a positive shock of the supply function over the efficiency of the *OTC* market deserves a more thorough discussion. Following the same procedure as in the previous section, we now analyze in more detail how the shocks applied in period 4 affect the traded quantity and the pattern of closing prices in the *OTC* market from period 4 onward. At the end of the section, we compare IM-inefficiency to EM-inefficiency in the *OTC* market.

Table 6 reports – for the first period of the second phase (i.e., period 4), on average over the second phase (periods 4–6), and on average over the third phase (periods 7–9) – the percentage of experimental sessions in which the traded quantity q in the *OTC* market is lower, equal or higher than the competitive quantity q^* , both without shocks (first line) and under each of the four types of shock.

	Period 4			Periods 4–6			Periods 7–9		
	$q < q^*$	$q = q^*$	$q > q^*$	$q < q^*$	$q = q^*$	$q > q^*$	$q < q^*$	$q = q^*$	$q > q^*$
No Shock	33.3	27.8	38.9	35.2	29.6	35.2	63.0	20.4	16.7
D^-	33.3	16.7	50.0	33.3	33.3	33.3	66.7	22.2	11.1
D^+	16.7	50.0	33.3	38.9	44.4	16.7	77.8	5.6	16.7
S^-	33.3	50.0	16.7	33.3	44.4	22.2	44.4	27.8	27.8
S^+	16.7	16.7	66.7	11.1	33.3	55.6	61.1	22.2	16.7

Table 6: Traded quantity q vs competitive quantity q^* in the *OTC* treatment with shocks (in percentage over all periods).

We note that the treatments where the percentage of sessions with $q < q^*$ in period 4 is lower than in the baseline (i.e. no-shock) treatment are those characterized by an increase of the traded quantity (D^+ and S^+). However, the positive shock to the supply function S^+ consistently induces a lower percentage of sessions for which $q < q^*$ over all trading periods after the shock, i.e. both in the second phase (11.1% vs. 35.2% in periods 4–6) and in the last phase of the experiment (61.1% vs. 63% in periods 7–9). Furthermore, after S^+ , the percentage of sessions with $q > q^*$ is higher than in the baseline treatment only for periods 4–6 and it tends to align to the baseline treatment in periods 7–9. Hence, the following result can be stated:

Result 6 (Traded quantity after shocks in *OTC*). *The introduction of a shock in the *OTC* market decreases the probability that $q < q^*$, and increases the likelihood that $q \geq q^*$ in the long run, only if the shock shifts downwards the supply function.*

Table 6 also shows that the introduction of a positive shock of the demand function in the *OTC* market does not lead to any substantial difference – in terms of traded quantity – with respect to the no-shock treatment. To this phenomenon corresponds the fact that the *OTC* efficiency raises more significantly as a consequence of a S^+ shock than of a D^+ shock, as shown above in Table 5. The higher efficiency reached in the S^+ case can be explained by looking at the pattern of closing prices in the *OTC* market. In the previous section we pointed out that in the *OTC* market closing prices converge to a price below the competitive price p^* . The fact that a S^+ shock reduces the predicted competitive price p^* brings the value of p^* closer to the actual *OTC* closing prices, thereby raising the efficiency of this mechanism. Figure 4 confirms this intuition. It presents the pattern of closing prices in an *OTC* market before (periods 1–3) and after (periods 4–9) a S^+ shock. The dashed line represents p^* , which in period 4 decreases as a consequence of S^+ , while the continuous line represents the average closing price, which before the shock is below p^* .

The price pattern in Figure 4 shows that, when p^* decreases as a consequence of an S^+ shock, the distance between the average closing price and p^* almost disappears. This happens because the predicted p^* decreases much more than the average closing price. In fact, buyers who are

Figure 4: Closing price patterns in an *OTC* treatment with shock S^+ . Dashed line: competitive price; continuous line: average closing price.

extra-marginal before the shock do not immediately realize (due to lack of public information in *OTC* about bids and asks) that they have become intra-marginal after the shock. We find that they accept closing prices slightly higher than the new equilibrium price. This corresponds to the above mentioned fact that, after the shock, the percentage of experimental sessions in which $q < q^*$ significantly decreases and the efficiency of the *OTC* mechanism increases. This effect cannot be produced by any of the other three types of shock. In fact, both S^- and D^+ increases p^* , thereby amplifying the gap between the average closing price and p^* . A D^- shock decreases p^* , but it also decreases $\bar{p}$ of a similar amount. In fact, sellers who are intra-marginal before the shock, do not immediately realize (due to lack of public information about bids and asks in *OTC*) that they are extra-marginal after the shock. We find that as soon as they realize they are not able to get rid of the unit of the good they own (same effect found in *OTC* with no shock) they ask for prices even lower than the average closing price before the shock. This is stated in the following result:

Result 7 (Closing price after shocks in *OTC*). *A positive shock of the supply function S^+ decreases the competitive price p^* more than the average closing price $\bar{p}$. Given Result 3, this reduces the (positive) difference between p^* and $\bar{p}$. This effect is not found under any of the other three types of shock.*

The above result can also be stated in terms of IM-inefficiency. In the previous section we showed that the inefficiency of the *OTC* mechanism is also due the failure of intra-marginal traders to exchange among themselves (IM-inefficiency), while this source of inefficiency is absent in the *DA* market. We also saw that IM-inefficiency occurs for sure when $q < q^*$. Since an S^+ shock drastically reduces the cases in which $q < q^*$ in *OTC* markets (Table 6), we conclude that this type of shock also reduces the IM-inefficiency of the *OTC* mechanism, thereby raising its

overall efficiency. This is confirmed by the regression analysis of efficiency and IM-inefficiency by treatment (see Table 9 in the Appendix): In the *OTC* treatment, the dummy variable for the S^+ shock type has a significant positive impact over the market efficiency ($P\text{-value} = 0.034$) and a negative impact ($P\text{-value} = 0.088$) over the IM-inefficiency. Notice that in the *OTC* treatment none of the other three shocks has a significant impact either over the market efficiency or the sources of market inefficiency.

Our regression analysis seems to point out also that, in the *DA* treatment, shocks to the supply curve tend to increase market efficiency from the baseline level of efficiency. On the contrary, shocks to the demand curve leave the market completely unaffected. Along the same lines, the tradeoff between IM and EM inefficiencies does not change after the market has suffered a shock. The latter result is not surprising, since the institutional framework of the *DA* market allows it to absorb shocks quite rapidly.

We can now state the last result of our research:

Result 8 (Efficiency after shocks in *DA* and in *OTC*). *The only shock that significantly affects efficiency in the *OTC* market in the long run is a positive shock of the supply function S^+ . Efficiency significantly increases with respect to the no-shock treatment thanks to a reduction in intra-marginal inefficiency. In the same fashion, only shocks to the supply function affect efficiency in the *DA* market, although they do not have any impact on the sources of market inefficiency.*

4 Summary and conclusions

Experimental economists have investigated the functioning and equilibrium properties of different market institutions, focusing on double auctions and other auction mechanisms. One key feature of auction markets is that buyers and sellers post their bids and asks publicly. However in many markets, such as those where real estate, currencies, or bulk commodities are traded, negotiations and transactions occur on a private and typically bilateral basis rather than through publicly posted bids and asks. Moreover, in an increasing number of these markets, bilateral negotiations and transactions take place via computer rather than orally. In order to clarify how this type of markets work, we have designed an *OTC* mechanism in which each agent looks for the best counterpart through private bids and asks submitted by computer. In a series of classroom experiments without monetary rewards that involved more than 3300 undergraduate students, we have studied the features and performance of this electronic *OTC* mechanism by taking as a benchmark the standard electronic *DA* mechanism. Recently, some theoretical models of *OTC* markets have been proposed (see, e.g., Duffie et al., 2005, 2007). However, to the best of our knowledge, our paper is the first that investigates electronic *OTC* markets from an experimental perspective.

We found that the loss of public information that characterizes our *OTC* market with respect to a *DA* market reduces the efficiency of the *OTC* mechanism by almost 8 efficiency point (Result 1). We also showed that this loss of efficiency is associated with two facts. First, in more than half of the trading periods, the quantity actually traded in *OTC* markets is lower than the competitive quantity. As subjects become more familiar with the *OTC* mechanism, actually

traded quantities increase, but this learning effect weakens when the trading period becomes shorter (Result 2). Second, in the *OTC* mechanism the average price at which the commodity units are traded is significantly lower than the competitive price (Result 3). We then discovered that, while the only source of inefficiency in *DA* markets is extra-marginal inefficiency (i.e., extra-marginal traders who exchange with intra-marginal traders), in *OTC* markets, inefficiency is also of the intra-marginal type. That is, there are some intra-marginal traders who could not exchange their good because of the lack of public information characterizing the *OTC* mechanism (Result 4). In the second part of the paper we introduced shocks into the picture, i.e., shifts in either the demand or the supply curve that modify the competitive equilibrium, and studied how efficiency is affected by different types of shocks. We found that, in the period when the shock takes place, the efficiency gap between the *DA* and the *OTC* mechanism shrinks. The latter result is independent of the type of shock. However, the gap emerges again after few periods and increases as the length of the trading period decreases (Result 5). Finally, we discovered that the only shock that significantly affects the functioning and efficiency of the *OTC* mechanism even in the long run is a positive supply shock, that is, a downward shift of the supply function that increases the competitive quantity and decreases the competitive price. This type of shock, in fact, decreases the cases in which the quantity actually traded in *OTC* markets is lower than the competitive quantity (Result 6); it reduces the difference between the competitive price and the average closing price in *OTC* markets (Result 7); and it also reduces the intra-marginal inefficiency associated with the *OTC* mechanism and thus increases its efficiency (Result 8).

From a policy perspective, these results suggest that regulators of financial markets and other markets in which interactions take place via computer and negotiations occur on a bilateral basis should prefer *DA* allocation mechanisms over *OTC* mechanisms, since the former warrants higher information disclosure and thus market equilibria that are more consistent with the competitive one. If an *OTC* mechanism is implemented, regulators should give traders more time to learn how the mechanism works than it would be the case in a *DA* market. Regulators should also give traders in an *OTC* market a sufficiently extended time to negotiate with other traders and find out the ‘best’ counterpart. Our experiment indicates, in fact, that too short transaction periods affect negatively and in a significant way the efficiency of *OTC* markets. Also in presence of unexpected modifications in the economic environment, which in our setting take the form of exogenous shocks, the *DA* mechanism seems to perform better than the *OTC* mechanism. Finally, if regulators are mainly concerned with enforcing a market mechanism that reveals a price very close to the competitive equilibrium, the *DA* market ought to be preferred. In particular, this work shows that prices in the *OTC* market are generally lower than the market equilibrium price, mainly due to persistent selling pressure: when public information about existing bids and asks is not available, sellers feel much more pressure than buyers in finding a trading counterpart. This in turn leaves room for positive surplus to buyers who should be excluded from transactions if the perfectly competitive equilibrium were performed. Hence, if an *OTC* mechanism is implemented, regulators should protect more (and eventually compensate) intra-marginal sellers, particularly those having costs closer to the equilibrium price.

These policy suggestions are admittedly very tentative, also because our paper seems to be the first systematic investigation of an electronic *OTC* mechanism from an experimental perspective. There is in fact ample room for further studies that explore experimentally the properties of this economically relevant market institution. In the companion paper Attanasi, Centorrino and

Moscatti (2014), we carry on the analysis of the *OTC* mechanism by integrating experimental and computational techniques and bringing into play zero-intelligence agents, that is, computer automata that post bids and asks at random. In future research, our *OTC* experimental design might be applied to the study of markets, such as those discussed by Rust and Hall (2003), in which agents trade through intermediaries and thus do not have access to public information about the bids and asks existing in the market.

References

- Ang, A., Shtauber, A. A. and Tetlock, P. C. (2013), ‘Asset Pricing in the Dark: The Cross Section of OTC Stocks’, *Review of Financial Studies* **26**, 2985–3028.
- Attanasi, G., Centorrino, S. and Moscati, I. (2014), Zero-Intelligence and Human Agents in an Experimental Over-the-Counter Market. Presented at the 2014 LABSI Workshop, University of Siena.
- Bardsley, N., Cubitt, R., Loomes, G., Moffatt, P., Starmer, C. and Sugden, R. (2009), *Experimental Economics: Rethinking the Rules*, Princeton University Press, Princeton, NJ.
- Blundell, R. and Bond, S. (1998), ‘Initial conditions and moment restrictions in dynamic panel data models’, *Journal of Econometrics* **87**(1), 115–143.
- Cason, T. N. and Friedman, D. (1996), ‘Price Formation in Double Auction Markets’, *Journal of Economic Dynamics and Control* **20**, 1307–1337.
- Cason, T. N. and Friedman, D. (2008), A Comparison of Market Institutions, in C. R. Plott and V. L. Smith, eds, *Handbook of Experimental Economics Results*, Vol. 1, Amsterdam - New York: North Holland, pp. 264–272.
- Chamberlin, E. H. (1948), ‘An Experimental Imperfect Market’, *The Journal of Political Economy* **56**(2), 95–108.
- Davis, D. D., Harrison, G. W. and Williams, A. W. (1993), ‘Convergence to Nonstationary Competitive Equilibria: An Experimental Analysis’, *Journal of Economic Behavior & Organization* **22**(3), 305–326.
- Duffie, D. (2012), *Dark Markets: Asset Pricing and Information Transmission in Over-the-Counter Markets*, Princeton, NJ: Princeton University Press.
- Duffie, D., Gârleanu, N. and Pedersen, L. H. (2005), ‘Over-the-Counter Markets’, *Econometrica* **73**(6), 1815–1847.
- Duffie, D., Gârleanu, N. and Pedersen, L. H. (2007), ‘Valuation in Over-the-Counter Markets’, *Review of Financial Studies* **20**(6), 1865–1900.
- Feldhütter, P. (2012), ‘The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures’, *Review of Financial Studies* **25**(4), 1155–1206.
- Fischbacher, U. (2007), ‘z-Tree: Zurich Toolbox for Ready-Made Economic Experiments’, *Experimental Economics* **10**(2), 171–178.
- Friedman, D. and Rust, J., eds (1993), *The Double Auction Market: Institutions, Theories, and Evidence*, Reading, MA: Addison-Wesley.
- Gode, D. K. and Sunder, S. (1993), ‘Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality’, *The Journal of Political Economy* **101**(1), 119–137.

- Guala, F. (2005), *The Methodology of Experimental Economics*, Cambridge University Press, Cambridge.
- Holt, C. A. (1996), ‘Classroom Games: Trading in a Pit Market’, *The Journal of Economic Perspectives* **10**(1), 193–203.
- Holt, C. A. (1999), ‘Teaching economics with classroom experiments’, *Southern Economic Journal* **65**(3), 603–610.
- List, J. A. (2002), ‘Testing Neoclassical Competitive Market Theory in the Field’, *Proceedings of the National Academy of Sciences* **99**(24), 15827–15830.
- List, J. A. (2004), ‘Testing Neoclassical Competitive Theory in Multilateral Decentralized Markets’, *Journal of Political Economy* **112**(5), 1131–1156.
- Ospina, R. and Ferrari, S. L. (2012), ‘A general class of zero-or-one inflated beta regression models’, *Computational Statistics & Data Analysis* **56**(6), 1609–1623.
- Rust, J. and Hall, G. (2003), ‘Middlemen versus Market Makers: A Theory of Competitive Exchange’, *The Journal of Political Economy* **111**(2), 353–403.
- Rust, J., Miller, J. and Palmer, R. (1993), Behavior of Trading Automata in a Computerized Double Auction Market, in Friedman and Rust (1993), pp. 155–198.
- Smith, V. L. (1962), ‘An Experimental Study of Competitive Market Behavior’, *The Journal of Political Economy* **70**(2), 111–137.

Appendix

	Period 4	Periods 4–6	Periods 7–9	Periods 1–9
No Shock	231.5 (0.028)	1959.5 (0.002)	2376.0 (0.000)	20084.5 (0.000)
D^-	16.0 (0.809)	220.5 (0.065)	248.0 (0.007)	2169.0 (0.000)
D^+	21.0 (0.686)	233.0 (0.025)	268.0 (0.001)	2249.0 (0.000)
S^-	22.0 (0.574)	247.0 (0.007)	291.0 (0.000)	2463.5 (0.000)
S^+	20.5 (0.739)	235.5 (0.018)	263.0 (0.001)	2029.5 (0.000)

Table 7: Mann Whitney U-statistics for the difference in efficiency (P -values in brackets).

	DA	OTC	$Joint$
Lag price	0.355 (0.000)	0.315 (0.000)	0.335 (0.000)
Buyer V	0.278 (0.000)	0.256 (0.000)	0.267 (0.000)
Seller C	0.315 (0.000)	0.381 (0.000)	0.346 (0.000)
Treatment			-0.691 (0.000)
Sargan test	0.000	0.000	0.000
AB test 1	0.000	0.000	0.000
AB test 2	0.000	0.001	0.000
Wald test	0.000	0.000	0.000

Table 8: Dynamic Panel Data regressions of closing prices over all periods (P -values in brackets).

	DA		OTC	
	Efficiency	IM-inefficiency	Efficiency	IM-inefficiency
Intercept	2.409 (0.000)	-1.950 (0.001)	2.285 (0.000)	-0.699 (0.000)
$(q - q^*)$	0.154 (0.000)	-0.661 (0.000)	0.250 (0.000)	-0.529 (0.000)
$ \bar{p} - p^* $	-0.038 (0.052)	-0.076 (0.115)	0.010 (0.453)	0.000 (0.980)
D^-	-0.145 (0.226)	0.322 (0.326)	0.018 (0.852)	-0.024 (0.858)
D^+	0.080 (0.517)	0.296 (0.536)	0.079 (0.417)	-0.062 (0.643)
S^-	0.177 (0.156)	0.399 (0.267)	0.085 (0.385)	-0.049 (0.720)
S^+	0.120 (0.339)	0.181 (0.569)	0.107 (0.288)	-0.267 (0.064)

Table 9: Beta regressions of efficiency and intra marginal (IM) inefficiency by treatment over all trading periods (*P-values* in brackets). Control variables for market period not reported.