

An Experimental Study of an Over-the-Counter Market

Giuseppe Attanasi^a, Samuele Centorrino^b, Ivan Moscati^c

^a*BETA, University of Strasbourg, and LERNA, Toulouse School of Economics, 61 Avenue de la Forêt-Noire, 67085 Strasbourg, France.*

Email: *attanasi@unistra.fr*

^b*Corresponding author. Toulouse School of Economics and GREMAQ, Manufacture des Tabacs, 21 Allée de Brienne, 31000 Toulouse, France.*

Current address: *Brown University, Department of Economics, 64 Waterman St, Providence, RI - 02912, USA.*

Phone: *+1-(401)863-5399.*

Email: *samuele.centorrino@tse-fr.eu*

^c*University of Insubria, Department of Economics, Via Monte Generoso 71, 21100 Varese, Italy.*

Email: *ivan.moscati@uninsubria.it*

Abstract

We study experimentally an over-the-counter (*OTC*) market mechanism that distinguishes itself from the standard double-auction (*DA*) mechanism by an important informational feature: in the *DA* setting traders post their bids and asks publicly, while in our *OTC* market each agent looks for the best counterpart through bilateral and private offers, and only the prices of closed transactions are made public. Although in many actual markets negotiations and transactions occur on a private, bilateral basis, *OTC* markets have not been previously studied in the experimental literature. We find that the lack of public information that characterizes our *OTC* mechanism with respect to a *DA* mechanism induces an efficiency loss. 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. This situation generates the so-called intra-marginal inefficiency. 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 shocks increasing the competitive quantity improve the *OTC*'s efficiency, and we clarify the causes of this phenomenon.

JEL classification: C92, D41, D47, D83.

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

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 reviews of this literature, see Holt, 1996 and 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 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 bids and asks 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, i.e. by making/receiving one offer at a time to/from a single counterpart. If the offer is accepted, the transaction is carried out, and the closing price is made public, a rule that may be labelled as *post-trade price transparency*. However, and differently from the *DA* setting, the history of asks and bids that led to that transaction remains private information. Also if the offer is withdrawn, and again in contrast to the *DA* setting, the offer 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 that investigated 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. However, there are three main differences

between Chamberlin's setting and ours. First, Chamberlin did not always make public the price of closed transactions, while we always implement post-trade price transparency. Second, 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. Third, in Chamberlin's classroom the bargaining between a buyer and a seller was conducted orally, so that 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.

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. Many types of government and corporate bonds, real estate, currencies, and bulk commodities are typically traded over the counter. In a number of these *OTC* markets, such as those for US 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, Ang et al., 2012, Duffie, 2012). But despite the economic relevance of *OTC* markets, we are not aware of any other experimental studies that have explored 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, we ran a series of classroom experiments involving more than 2400 undergraduate students of almost the same age over a period of four years (from 2009 to 2012) 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.

Our (natural) 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 asks and bids is public information, makes *OTC* markets less efficient than *DA* markets. Our experimental data 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 potential surplus. We find that in *DA* markets the average efficiency index is 95.5%, while in *OTC* markets the efficiency index is 86.7%. Thus the information gap between the *OTC* and the *DA* settings determines a loss of efficiency of almost 9%.

To better understand how the lack of information about the history of asks and bids affects negatively the efficiency of the *OTC* mechanism, we study the pattern of trading prices and traded quantity in both the *OTC* and the *DA* setting. We find that, because of its informational features, in the *OTC* mechanism trading 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, shocks that decrease the competitive quantity do not affect substantially the efficiency of the *OTC* mechanism, while shocks that increase the competitive quantity significantly raise the efficiency of the *OTC* markets. In the final part of the paper we clarify the causes of this phenomenon.

2. Experimental Design

Our experiments rely on a design that follows Smith (1962), in that valuations and costs are exogenously given, and Cason and Friedman (1996), in that subjects are allowed to trade only one unit per period. We have two main treatments, *DA* and *OTC*. In both of them subjects are randomly assigned the role of either buyer or seller. A single homogeneous good can be bought or sold during a given number of periods of finite length. Each buyer (seller) is assigned a redemption value for the single unit of the good he/she has to buy (sell). 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. In both treatments a feasibility constraint is imposed: buyers cannot bid over their own valuation, and sellers cannot ask under their own cost.

Two main features distinguish the *DA* from the *OTC* treatment. First, as already mentioned, in the *DA* setting buyers and sellers post their bids and asks publicly, so that the bid-ask history of the market is common knowledge. In the *OTC* setting, by contrast, each agent looks for the best counterpart through private bids and asks, and only the prices of closed transactions are made public. Second, in the *DA* mechanism the *bid/ask improvement rule* holds. In a *DA* market every agent is informed about the highest bid and the lowest ask existing in the market. The bid/ask improvement rule imposes the stipulation that, to make a valid offer, an agent has to improve on the existing situation: a buyer has to submit a bid higher than the current highest bid, and a seller has to submit an ask lower than the current lowest ask. 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. With the *OTC* mechanism, however, the bid/ask improvement rule does not hold: the agents do not observe the best bid and ask present in the market, and can also replace their current bid (ask) with a lower (higher) one.

We ran computerized classroom experiments implementing this design through the z-Tree software (Fischbacher, 2007). Sessions were held at Bocconi University, Milan, during a first-year introductory course in Microeconomics over four consecutive years, from 2009 to 2012, always in October (first semester). About one third of enrolled students per year were involved in the experiments, i.e. 2416 students as a whole. The four 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 60 experimental sessions we ran.¹ The traders were divided equally into buyers and sellers (20 buyers and 20 sellers for each session). We ended up holding 30 sessions for the *DA* treatment with a total of 1204 subjects; and 30 sessions for

¹To be precise, in 48 sessions we had 40 subjects, in 10 sessions we had 42 subjects, and in 2 sessions we had 38 subjects.

the *OTC* treatment with 1212 subjects in total. All classroom experiments were run by the same experimentalist (G. Attanasi), who is also one of the authors of this paper.

Each experimental session consists of six market periods of equal clock time length, namely 120 seconds per period. The six periods are partitioned into two phases, with each phase consisting of three periods and lasting 360 seconds each, that is, 6 minutes per phase. Subjects are informed about the existence of a second phase only at the end of the first phase.²

At the beginning of each phase, subjects are given three pieces of information: their ID, their role (either buyer or seller) and their redemption value. Neither in the *DA* treatment nor in the *OTC* treatment do subjects know about the distributions of valuations and costs in the market. While their role and redemption value remain equal until the end of the phase, their ID changes at every period. This prevents subjects from identifying trading counterparts in a given period on the basis of IDs learned in previous periods. Every seller owns only one unit of the good and his/her cost sets the minimum amount he/she can cash in for that unit. The buyer's valuation sets instead the maximum amount he/she can spend for one unit of the good. The subject's redemption value is private information and it always appears on his/her screen. At the end of the first phase, roles are kept constant while redemption values are reshuffled. Therefore, in the second phase a given subject might have a different redemption value.

In the *DA* treatment, bids and asks are public information, every buyer (seller) is always informed about the best bid (ask) on the market, and the bid/ask improvement rule holds. Once an offer is accepted, the closing price appears on every subject's screen. In the *OTC* treatment, a subject can send an offer to a single counterpart, by indicating the amount of the offer and the counterpart's ID. Only one offer can be submitted at a time. If the offer is accepted, the closing price appears on every subject's screen. If the counterpart does not reply, the offer may be withdrawn and a new offer can be made that differs either in terms of the amount, the counterpart's ID or both.³ However, bids and asks submitted to a specific subject are not publicly disclosed, and no bid/ask improvement rule holds.

At the end of each period, and for both treatments, each subject sees on the screen his/her payoff as the difference between valuation and closing price - if he/she is buyer - or between closing price and cost - if he/she is seller. If a subject does not trade his/her commodity unit within the period, his/her payoff is equal to zero. Being in a classroom experiment, subjects are not remunerated for their participation in the experiment. However, we give them an incentive to play fairly: at the end of each of the two phases, subjects are ranked according to their total profit in that phase.⁴ We then ask the four subjects having earned the highest and lowest total profit in

²In order to rule out the possibility that our experimental outcomes might be modified by further learning on the part of the subjects, we ran a shorter third phase still consisting of three periods but lasting only 3 minutes in total (again, this third phase was announced only at the end of the previous one). The outcomes of the third phase, which are available upon request, are substantially identical to those of the second phase. This makes us confident that our results do not change when subjects gain more experience, and allows us to focus on the first two phases only.

³If a subject receives more than one offer at a time, those offers are automatically ranked, so that the best deal always appears on the top of his/her screen.

⁴Profits 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.

that phase to stand up. The former are praised publicly for their performance. The latter instead may be flouted by classmates.

In order to get deeper insights about the *OTC* mechanism, we introduce shocks into the picture and study how efficiency in the *DA* and the *OTC* mechanisms is affected by shocks. Subjects are not given any information about shocks, which are defined as shifts in either the demand or the supply curve that lead to a change in the competitive quantity q^* and the competitive price p^* . In 40 experimental sessions (that is, 2/3 of the overall sessions) we apply shocks to either demand or supply for both the *DA* treatment (20 sessions) and the *OTC* treatment (20 sessions). In all these sessions, a shock occurs in period 4 and is maintained during periods 4 – 6.

We implement four different types of shocks (5 sessions for each type of shock, per treatment):

1. a positive (upward) shift of demand, indicated as D^+ and leading to an increase of both q^* and p^* ;
2. a negative (downward) shift of demand, indicated as D^- and leading to a decrease of both q^* and p^* ;
3. a positive (downward) shift of supply, symbolized by S^+ and determining an increase of q^* and a decrease of p^* ;
4. a negative (upward) shift of supply, symbolized by S^- and determining a decrease of q^* and an increase of p^* .

Table 1 summarizes our experimental design by indicating the absolute number of sessions we ran for each type of treatment (*DA* and *OTC*), without shock and for each type of shock (D^+ , D^- , S^+ , S^-).

	No Shock	D^+	D^-	S^+	S^-
DA	10	5	5	5	5
OTC	10	5	5	5	5

3. Results

The basic measure we use to compare the performance of *DA* and *OTC* markets is the efficiency index as defined in Gode and Sunder (1993), that is, the ratio between the total surplus actually realized from trade and potential surplus. The efficiency index goes from 0 (minimal efficiency) to 100 (full efficiency), and the higher the index value the more efficient the market. The efficiency index, however, says nothing about the causes of the inefficiency. To better understand these causes, we study the pattern of trading prices in the *DA* and *OTC* markets, and compare the quantity actually traded in these markets with the competitive quantity. Moreover, following Rust et al. (1993), we decompose inefficiency into two main components: intra-marginal inefficiency and extra-marginal inefficiency, both of which we will illustrate shortly.

We compare the relative performance of *DA* and *OTC* markets starting from the treatment without shocks, and we then introduce shocks, which are applied at the beginning of the fourth

market period, that is, in the first period of phase 2. In order to facilitate comparison between the treatments without shocks and with shocks, we report results for the entire experimental session (periods 1 – 6) and for periods 1 – 3 and 4 – 6, separately.⁵

3.1. OTC and DA without shocks

Table 2 reports the efficiency index for the *DA* and *OTC* markets respectively. The informational advantage of the *DA* over the *OTC* mechanism is due to the fact that in the former the entire history of asks and bids 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 – 6 we observe an average efficiency index of 95.5 in the *DA* market, and an average efficiency index of 86.7 in the *OTC* market.⁶ Thus the informational disadvantage of the *OTC* mechanism compared with the *DA* mechanism determines a significant loss of efficiency of almost 9%.

Table 2: Efficiency index, with no shocks

	Periods		
	1 – 6	1 – 3	4 – 6
DA	95.5	94.7	96.3
OTC	86.7	83.3	90.1

As 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. This is true, in particular, for the *OTC* market, where efficiency increases by almost 7% in passing from periods 1 – 3 to periods 4 – 6. This result is interesting by itself: in the *OTC* market, learning partially offsets the relative lack of information.

Higher inefficiencies in *OTC* markets may come from two sources: the actually traded quantity q can be different from the competitive quantity q^* ; and/or trading prices may converge to a price that is different from the competitive price p^* .

Table 3 reports the difference between q and q^* . The study of this difference gives us important information about the causes of the inefficiency of a market institution. If $q < q^*$, then profitable trades between some intra-marginal buyer, i.e. a buyer with a valuation higher than the competitive price, and some intra-marginal seller, i.e. a seller with a cost lower than the competitive price, have not taken place. If $q > q^*$, then some units that should have remained out of the market have instead been exchanged: either some extra-marginal buyer, i.e. a buyer with a valuation lower than the competitive price, managed to buy from an intra-marginal seller, or some extra-marginal seller, i.e. a seller whose cost is higher than the competitive price, 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.

In the *DA* market, $q < q^*$ in only 12% of all market periods. This result is consistent with existing experimental evidence about *DA* (see, e.g. Gode and Sunder, 1993, Cason and Friedman,

⁵All data and instructions are available upon request.

⁶The results we have obtained for the *DA* market are in line with the findings of Gode and Sunder (1993), and Cason and Friedman (1996).

Table 3: Traded quantity q vs competitive quantity q^* , with no shocks.

	Periods 1 – 6			Periods 1 – 3			Periods 4 – 6		
	$q < q^*$	$q = q^*$	$q > q^*$	$q < q^*$	$q = q^*$	$q > q^*$	$q < q^*$	$q = q^*$	$q > q^*$
DA	12%	39%	49%	9%	54%	37%	15%	24%	61%
OTC	70%	18%	12%	76%	12%	12%	64%	24%	12%

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 easier 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 deals that can be accepted by an intra-marginal counterpart. Thus, the *DA* mechanism facilitates profitable exchanges between intra-marginal traders, and accordingly only rarely (12% of the time) delivers a traded quantity lower than the competitive quantity.

By contrast, in the *OTC* market $q < q^*$ in 70% of all market periods. Our explanation for this result is that, because in *OTC* markets negotiations are conducted on a one-to-one basis and market time is limited, it is easier for intra-marginal traders to miss the possibility of closing a profitable transaction. To reinforce this explanation, notice that, when traders become more experienced, the percentage of periods for which $q < q^*$ decreases: in periods 1 – 3, it is 76% on average, while in periods 4 – 6, it is 64% on average. This corresponds to an increase in efficiency of the *OTC* mechanism, as reported in Table 2.

The result that in the *OTC* market $q < q^*$ can be related to the pattern of closing prices. If the market converges to an equilibrium where the average price is 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 their convergence to the equilibrium⁷. In the left panel of Figure 1, we draw the pattern of closing prices over periods 1 – 6 for a single *DA* experimental session, while on the right panel we do the same for a single *OTC* session. In both panels, the units traded are plotted in abscissas according to the order in which they have been traded (the first unit traded comes first on the abscissas line, the second unit comes second, etc.), while the prices corresponding to each traded unit are plotted in ordinates. In both panels, the dashed line corresponds to the competitive price, while the continuous line expresses the average closing price in, respectively, the *DA* and *OTC* experimental sessions.

In the first place, we notice that price convergence 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.

This pattern of closing prices does not characterize only the specific *DA* and *OTC* experimental

⁷For the sake of clarity and simplicity, we focus on a graphical analysis of convergence. More detailed statistical tests are available upon request.

Figure 1: Closing price patterns in a *DA* market (left panel) and an *OTC* market (right panel) when there is no shock on the market. The dashed line indicates the competitive price. The continuous line indicates the average closing price.




sessions represented in Figure 1, but holds for all the *DA* and *OTC* experimental sessions over the last three periods, when price convergence is most likely to occur. This is shown in Figure 2, where we report the histogram of the relative deviations from the competitive price *for periods 4-6 of all experimental sessions*. We plot the histogram of $\Delta p_i^* = (p_i - p^*)/p^*$, where p_i is the closing price of unit i and p^* is the competitive price only for the last three periods.

In Figure 2(a), 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, the closing prices converge to p^* . In contrast, the histogram for the *OTC* market - Figure 2(b) - shows that relative deviations of closing prices are almost normally distributed and that their mean is slightly below 0. This confirms 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.

We can bring the analysis even further. To understand and compare better the sources of the inefficiencies in the *DA* and *OTC* markets, it is useful to decompose the loss of surplus generated by these two market mechanisms into *intra-marginal inefficiency* and *extra-marginal inefficiency*, two notions introduced by Rust et al. (1993). There is intra-marginal inefficiency (henceforth *IM-inefficiency*) when two intra-marginal traders do not exchange. Extra-marginal inefficiency (*EM-inefficiency*), by contrast, 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 decreases the traded quantity: if two intra-marginal traders do no exchange, q is lower than q^* . EM-inefficiency, by contrast, either does not modify q or increases it. To see why, recall that exchange between two extra-marginal traders is impossible. Thus, an extra-marginal trader always exchanges his/her commodity unit with an intra-marginal trader, and in so doing he/she displaces some intra-marginal trader and the latter's commodity unit. Two things may happen to a displaced intra-marginal trader: he/she

Figure 2: Histograms of the absolute relative deviation of closing prices from the equilibrium price, periods 4 – 6 in *DA* (left panel) and *OTC* (right panel).



may be unable to trade his/her unit, in which case the quantity q traded on the market does not change; otherwise, he/she may find another extra-marginal trader with which he/she can exchange his/her unit and in this latter case q increases. 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 lower than q^* .

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

1. if $q > q^*$, the only source of inefficiency is EM-inefficiency;
2. if $q = q^*$ and actual surplus equals potential surplus, the inefficiency equals 0: all intra-marginal traders have traded their unit and all extra-marginal traders are out of the market;
3. if $q = q^*$ but actual surplus is lower than potential surplus, then the existing efficiency is certainly due to EM-inefficiency. IM-inefficiency is ruled out because q is not smaller than q^* ;
4. 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 price and quantity, the potential surplus, the closing prices of all traded units, and the quantity actually traded - we can decompose the loss of surplus associated with the *DA* and *OTC* markets into their IM-inefficiency and EM-inefficiency components. Table 4 reports the results of our inefficiency audit.

We find that in the *DA* market IM-inefficiency accounts for only 8% of total inefficiency, while EM-inefficiency accounts for the residual 92% percent. This is congruous with the data displayed in Table 2, which show that in *DA* markets the traded quantity q is lower than the competitive quantity q^* only 12% of times (remember that we have IM-inefficiency, possibly associated with

Table 4: Sources of inefficiency by treatment with no shocks over periods 1-6 (in percentage over total inefficiency).

	Periods 1 – 6	
	EM-inefficiency	IM-inefficiency
DA	92	8
OTC	54.2	45.8

EM-inefficiency, only when $q < q^*$). In the *OTC* setting, by contrast, we have a mixture of IM- and EM-inefficiency: on average across all sessions, IM-inefficiency accounts for about 46% of the total market inefficiency. Again, this finding is consistent with the data of Table 3 according to which in *OTC* markets q is lower than q^* 70% of the time.

Result 1. *The OTC mechanism is, on average, less efficient than the DA mechanism. The efficiency loss due to informational disadvantage of the OTC mechanism can be quantified at around 9%.*

This efficiency loss is due to the fact that in the OTC mechanism closing prices converge to a price below the competitive price p^ and that, accordingly, the traded quantity q is lower than the competitive quantity q^* . This situation leaves some intra-marginal traders out of the market and generates a significant level of IM-inefficiency.*

3.2. Introducing shocks

In order to grasp better the functioning of the *OTC* mechanism, we introduce shocks into the picture and study how the efficiency of the *OTC* and the *DA* mechanisms 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^* . They occur in period 4 and are maintained during periods 4 – 6. We implement four types of shock: D^+ , which increases q^* and p^* ; D^- , which decreases q^* and p^* ; S^+ , which increases q^* but decreases p^* ; and S^- , which decreases q^* and increases p^* .

Table 5 reports the efficiency index for *DA* and *OTC* markets in period 4 and over periods 4 – 6 in the case without shocks (first column) and for the four types of shock.⁸

Table 5: Efficiency index by treatment and shock.

	No Shock		D^+		D^-		S^+		S^-	
	Periods		Periods		Periods		Periods		Periods	
	4	4 – 6	4	4 – 6	4	4 – 6	4	4 – 6	4	4 – 6
DA	97.3	96.3	97.2	97.6	92.1	94.7	97.2	98.1	89.4	92.3
OTC	90.3	90.1	97	94.3	91	87.9	98.5	94.5	86.7	88.7

Consider first the *DA* market. Without shocks, the efficiency index in period 4 is equal to 97.3. When shocks increasing q^* are implemented (i.e. D^+ and S^+) efficiency stays roughly constant in period 4: for both D^+ and S^+ the efficiency index is 97.2. However, for shocks decreasing q^* (i.e. D^- and S^-) in period 4 efficiency decreases: it falls to 92.1 for D^- and even more sharply (89.4)

⁸Although 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.

for S^- . Over periods 4 – 6, for D^+ and S^+ efficiency remains high, while for D^- and S^- it tends to increase and return to pre-shocks levels.

These findings are 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 the EM-inefficiency of the *DA* mechanism.

In the *OTC* market we observe a different impact of shocks. In period 4 shocks D^+ and S^+ sharply raise efficiency: from the baseline of 90.3, in period 4 efficiency reaches 97 under a D^+ shock, and 98.1 under a S^+ shock. By contrast, efficiency stays almost constant under a D^- shock (91), and decreases under a S^- shock (86.7). However, over periods 4 – 6 efficiency tends to return to pre-shocks levels also in the *OTC* market.

To explain the positive effect of shocks D^+ and S^+ on the efficiency of the *OTC* mechanism in period 4, we analyze in more detail how these shocks affect the traded quantity and the pattern of closing prices. Table 6 reports the fraction of sessions in which, in period 4, the traded quantity q in *OTC* markets is lower, equal or higher than the competitive quantity q^* , both without shocks (first line) and with the four shocks (lines 2-5).

We note that, in treatments D^+ and S^+ , the fraction of sessions in which $q < q^*$ drastically decreases. From a 50% in the control treatment, it drops to 20%, under D^+ , and to 0%, under S^+ . In the previous section we showed that the inefficiency of the *OTC* mechanism is mainly due the failure of intra-marginal traders to exchange among themselves, that is, to IM-inefficiency. We also saw that IM-inefficiency occurs only when $q < q^*$. Since shocks D^+ and S^+ drastically reduce the cases in which $q < q^*$, we conclude that these shocks also reduce the IM-inefficiency of the *OTC* mechanism, thereby raising its overall efficiency.

Table 6: Traded quantity q vs competitive quantity q^* in the *OTC* market with shocks.

	Period 4		
	$q < q^*$	$q = q^*$	$q > q^*$
No shock	50%	40%	10%
D^+	20%	40%	40%
D^-	40%	0%	60%
S^+	0%	20%	80%
S^-	80%	0%	20%

There is an additional aspect that should be clarified. The fraction of sessions in which $q < q^*$, in period 4, is significantly lower under a S^+ shock (0%) than under a D^+ shock (20%). To this phenomenon corresponds the fact that the *OTC* efficiency raises more significantly as a consequence of a S^+ shock (from 90.3% to 98.5%) than of a D^+ shock (from 90.3% to 97%).

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 p^* brings the value of p^* closer to the actual *OTC* trading prices and thus further raises the efficiency of this mechanism.

Figure 3 confirms this intuition. It presents the pattern of closing prices in an *OTC* market before and after 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 trading price, which before the shock is below p^* . We see that when p^* decreases as a consequence of the S^+ shock, the distance between the average trading price and p^* almost disappears. This corresponds to the fact that after the shock the fraction of cases in which q is smaller than q^* goes to zero and the efficiency of the *OTC* mechanism increases.

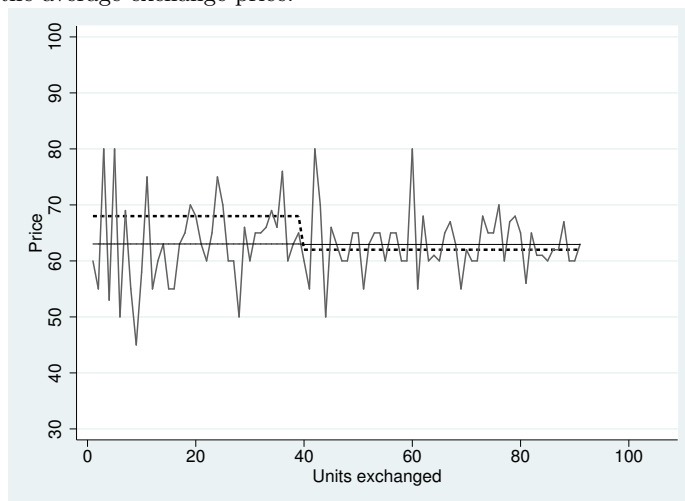
We can sum up our discussion of how the efficiency of the *OTC* mechanism is affected by shocks as follows:

Result 2. *In the short run, shocks that reduce the competitive quantity do not affect substantially the efficiency of the OTC mechanism, while they decrease the efficiency of the DA mechanism.*

In the short run, shocks that raise the competitive quantity do not affect substantially the efficiency of the DA mechanism, while they significantly increase the efficiency of the OTC mechanism. This positive effect on the OTC's efficiency is due to a drop in IM-inefficiency.

In the long run, for all shocks the efficiency of both mechanisms tends to return to pre-shocks level.

Figure 3: Closing price patterns in an *OTC* market with S^+ . The dashed line indicates the equilibrium price. The continuous line indicates the average exchange price.



4. Conclusions

Perfectly competitive markets remain a fundamental theoretical benchmark for economists. Many existing markets, and especially financial markets, display a number of features that characterize competition (many buyers and sellers, who are “small” with respect to the market and trade a homogeneous commodity) but are not perfectly competitive. For more than half a century, experimental economists have studied market institutions that have a perfect-competition character, and have compared the equilibrium and efficiency properties of these mechanisms to the competitive benchmark. In particular, experimental economists have focused on double auctions and other auction mechanisms.

One key feature of auction mechanisms is that buyers and sellers post their bids and asks publicly. However, in many “almost-competitive” 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. In order to clarify how these markets work, we have designed an over-the-counter mechanism where each agent looks for the best counterpart through private bids and asks. We have studied experimentally the features and performance of this mechanism by taking the double-auction mechanism as a benchmark. We found that the loss of public information that characterizes our *OTC* market with respect to a *DA* market alters the equilibrium properties of the *OTC* mechanism and reduces its efficiency by about 9%. From a policy perspective, this result suggests that regulators of financial markets should prefer *DA* allocation mechanisms over *OTC* mechanisms, for the former warrant higher information disclosure and thus market equilibria that are more consistent with the competitive one.

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 *OTC* markets from an experimental perspective. Thus there is room for further studies that explore experimentally the properties of this economically relevant market institution. In a companion paper (Attanasi et al., 2013), 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.

Acknowledgement

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, and participants in the 3rd Annual Xiamen University International Workshop on Experimental Economics and Finance, and in the 2011 Labsi Conference in Capua 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). The usual disclaimer applies.

References

- Ang, A., Shtauber, A. A. and Tetlock, P. C. (2012), Asset Pricing in the Dark: The Cross Section of OTC Stocks. Mimeo, Columbia University.
- Attanasi, G., Centorrino, S. and Moscati, I. (2013), Zero-Intelligence and Human Agents in an Experimental Over-the-Counter Market. Mimeo, Toulouse School of Economics.
- 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.
- 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.
- Holt, C. A. (1996), 'Classroom Games: Trading in a Pit Market', *The Journal of Economic Perspectives* **10**(1), 193–203.
- 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.