Zero-Intelligence vs. Human Agents: An Experimental Analysis of the Efficiency of Double Auctions and Over-the-Counter Markets of Varying Sizes.

Working Paper 2020-04

Giuseppe Attanasi, Samuele Centorrino and Elena Manzoni

March, 2020

Zero-Intelligence *vs.* Human Agents: An Experimental Analysis of the Efficiency of Double Auctions and Over-the-Counter Markets of Varying Sizes

Giuseppe Attanasi (University of Côte d'Azur, Nice)

Samuele Centorrino (Stony Brook University) **Elena Manzoni** (University of Verona)

March 26, 2020

Abstract

We study two well-known electronic markets: an over-the-counter (OTC) market, in which each agent looks for the best counterpart through bilateral negotiations, and a double auction (DA) market, in which traders post their quotes publicly. We focus on the DA-OTC efficiency gap and show how it varies with different market sizes (10, 20, 40, and 80 traders). We compare experimental results from a sample of 6,400 undergraduate students in Economics and Management with zero-intelligent (ZI) agent-based simulations. Simulations with ZI traders show that the traded quantity (with respect to the efficient one) increases with market size under both DA and OTC. Experimental results with human traders confirm the same tendency under DA, while the share of periods in which the traded quantity is higher (lower) than the efficient one decreases (increases) with market size under OTC, ultimately leading to a DA-OTC efficiency gap increasing with market size. We rationalize these results by putting forward a novel game-theoretical model of OTC market as a repeated bargaining procedure under incomplete information on buyers' valuations and sellers' costs, showing how efficiency decreases slightly with size due to two counteracting effects: acceptance rates in earlier periods decrease with size, and earlier offers increase, but not always enough to compensate the decrease in acceptance rates.

JEL codes: C70, C91, C92, D41, D47

Keywords: Market Design, Classroom Experiment, Agent-based Modelling, Game-theoretic Modelling.

Acknowledgements: We thank Dhananjay K. Gode and Shyam Sunder for sharing the code for agent-based simulations of zero-intelligent computerized trading in double auction markets. We thank Antonio Bisignano, Sara Gil Gallen, Stefano Manzoni, and Chris Ouangraoua for useful research assistantship. We are grateful to Olivier Armantier, Vardan Baghdasaryan, Ivan Moscati, Simone Quercia and seminar participants at the 2019 VELE meetings (University of Verona) and the 2019 IESC workshop in Cargèse for useful comments. The research leading to these results has received funding from the French Agence Nationale de la Recherche (ANR) under grant ANR-18-CE26-0018-01 (project GRICRIS).

1 Introduction

Experimental markets have been studied to understand the equilibrium properties and efficiency of different market structures. The most common market structure considered by the experimental literature is Vernon Smith's (1962) double auction (henceforth, DA) mechanism, which has been designed so far to test the behavior of competitive markets such as financial markets (Friedman and Rust, 1993; Plott, 2008; Cason and Friedman, 2008), and to prove that automata can do as well as humans when they trade under simple rules (see Gode and Sunder, 1993, 1997, 2004, 2018).

In a DA market, buyers and sellers typically 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. Given that different units of the commodity can be traded at different prices, and traders are price makers, DA markets are non-competitive markets. However, transaction prices and traded quantity quickly converge to the competitive price and quantity, and the efficiency reached by DA markets closely approximates that reached by competitive markets.

In the literature, also decentralized markets have been studied experimentally, with the most well-known example represented by over-the-counter (henceforth, OTC) markets (see, e.g., Chamberlin, 1948; List, 2002, 2004). In OTC markets traders individually look for their counterparts and trades happen through private bilateral negotiations. There exist many types of OTC markets, which differ in features such as the exact process through which each trader searches for a counterpart or the possible presence of intermediating traders such as brokers. There are, however, two main features characterizing all OTC markets that make them comparable to DA markets. First, as in DA markets, different buyers and sellers can trade the same commodity at different prices, therefore they are price makers: also OTC markets are non-competitive markets. Second, differently from DA markets, where pre-trade is public, under OTC agents' private bargaining gives little information about trading opportunities, i.e., bids and asks of the other traders.

Attanasi et al. (2016) have compared the relative performance of OTC to DA markets with 40 human traders. They impose public information about trading prices in OTC markets, a feature that holds by construction in DA markets. They experimentally find that market decentralization (private trading) determines a loss of efficiency of almost 8 efficiency points in OTC vs. DA markets. They associate this DA-OTC efficiency gap to the lack of pre-trade price transparency in OTC markets, and highlight the important role of information about the entire history of bids and asks that characterizes DA (centralized) markets.

As shown by Smith (1962) and subsequent experimental research (see, e.g., Attanasi et al., 2019), the convergence and efficiency properties of DA markets with human traders are robust to modifications of the market size, i.e., of the number of buyers and sellers (for a review, see Friedman and Rust, 1993). The same occurs under agent-based simulations with automata (for a review, see Gode and Sunder, 2018; Rust et al., 2018). To the best of our knowledge, similar tests of robustness to market size modifications are missing for OTC markets, in both above mentioned streams of research. Therefore, in this paper we investigate OTC markets through the simultaneous use of experiments with human traders and agent-based simulations with automata, with a threefold aim.

First, we want to understand whether humans may perform as well as automata in decentralized markets. In this regard, to the best of our knowledge, we are the first to extend Gode and Sunder (1993) agent-based simulations to OTC markets. Second, we aim to verify whether the efficiency loss of OTC markets with human agents with respect to DA markets with human agents varies with the market size. Third, we want to determine the trend of the (supposed) efficiency loss of OTC markets with human agents with respect to OTC markets with automata as a function of the market size. In this analysis, we consider the efficiency loss between humans and automata in DA markets as a control.

For what concerns this last aim, we also acknowledge an asymmetry in the theoretical literature on the effects of market size on efficiency. On the one hand, there is a vast literature on the effects of market size on the efficiency of DA markets with human agents. Although the theoretical results on the efficiency of DA markets are very sensitive to the institutional features of the market itself, they go in the direction of showing that, when the number of traders is small, both efficient and inefficient equilibria might coexist, while, as the number of traders grows, the trading outcome converges to the competitive equilibrium (Friedman, 2018). Theoretically, the analysis of efficiency of DA markets with one-way traders (sellers and buyers), independent values and single indivisible units has been developed by Chatterjee and Samuelson (1983), Satterthwaite and Williams (1989), Gresik and Satterthwaite (1989), Wilson (1985), Cripps and Swinkels (2006). All these papers highlight how a static DA market converges to 100% efficiency as the number of traders grows (at different rates depending on the specific institutional characteristics of the market). Wilson (1987) confirms that also continuous DA markets with a large number of traders tend to 100% efficiency. Specifically, Wilson (1987) argues that this is due to an increased competitive pressure among traders of the same type (buyers or sellers), who accept advantageous offers more quickly to avoid being anticipated in acceptance by other traders of the same type.

On the other hand, the theoretical literature on how efficiency of OTC markets is affected by market size is missing. Existing models of OTC markets (Duffie et al. (2007) and follow-up papers) do not focus on the strategic effects which may influence how the agents' behavior in OTC markets varies with size. For this reason we put forward a novel simple model of OTC markets which allows us to postulate experimental predictions on the effects of size on the efficiency of OTC markets with human agents. Our game-theoretical model shows that the intuition of Wilson (1987) on why DA markets with large number of traders become more efficient does not apply to OTC markets. As the offer is made to a single counterpart, increasing the number of traders does not increase the competitive pressure from traders of the same type, while it increases the probability of receiving more advantageous offers from traders of the opposite type, thereby leading traders to strategically wait more before accepting an advantageous offer. As a consequence, an increase in market size reduces efficiency in OTC markets.

We tested our theoretical predictions in a series of computerized classroom experiments (300 market sessions), by involving a gender balanced sample of 6,400 undergraduate students of the same age ($19-20$ years old), nationality (mostly Italians), and field of study (Economics and Management), during a first-year introductory course in Microeconomics at Bocconi University Milan over five consecutive academic years, namely from 2015 to 2019 (1,280 students per year on average). We implemented a 2X4 between-subject design, with each student participating in only one treatment, characterized by one out of 2 trading mechanisms (DA or OTC) and one out of 4 market sizes with equal number of buyers and sellers (10, 20, 40, or 80 traders). Due to the exceptionally large number of students involved, and as is indeed common in many classroom experiments (see, e.g., Holt, 1996, 1999; Attanasi et al., 2016), we did not use monetary incentives. Rather, we incentivized students to play effectively by publicly praising the best performing traders among them (more details in Section 2). Correspondingly, we ran 800 periods of agent-based simulations for each of the 8 treatments (6,400 periods in total) by implementing algorithms of zero-intelligent (ZI) computerized trading (automata) *à la* Gode and Sunder (1993) under each trading mechanism – market size combination. This was meant to disentangle the (automatized) effect of market rules from the (human) effects of e.g. learning and strategic behavior, thereby providing a control for both DA and OTC's efficiency under different market sizes.

Our results on the comparison between OTC markets and DA markets with different types of traders (human vs. automata) and different market sizes (10, 20, 40 and 80 traders) basically confirm our theoretical predictions. Referring to the three above mentioned research questions, first of all we find that, as for DA markets, in OTC markets human traders reach lower level of efficiency than ZI agents. Furthermore, the DA-OTC efficiency loss with human agents increases with market size. This is due to the fact that, while with ZI agents the traded quantity (with respect to the efficient one) increases with market size under both DA and OTC, with human agents we detect the same tendency under DA, while the share of periods in which the traded quantity is higher (lower) than the efficient one decreases (increases) with market size under OTC. The latter result is the key finding of our game-theoretical model of bargaining under OTC.

The rest of the paper is organized as follows. In Section 2 we illustrate the experimental design and the corresponding control via agent-based simulations. In Section 3 we present the novel theoretical model of OTC markets (Section 3.1) and the experimental predictions (Section 3.2). In Section 4 we present our experimental findings. In Section 5 we conclude and discuss some policy implications of our experimental findings.

2 Design

In this section we first describe the features of the experimental design with human agents (Section 2.1) and then we describe the features of the agent-based simulations with ZI agents (Section 2.2).

2.1 Experimental design with Human agents

Procedures. We ran computerized classroom experiments implementing the design through the z-Tree software (Fischbacher, 2007). All experiments were held at Bocconi University, Milan, during a first-year introductory course in Microeconomics over five consecutive years, from 2015 to 2019, always in the first two weeks of October (first semester), in the same computerized room, and administered by the same experimenter, who is also one of the authors of this paper (G. Attanasi). The five cohorts of participants were homogeneous in many relevant characteristics: 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 or Management).[1] Each experimental session was characterized by a market mechanism (DA or OTC) and a market size of $n \in \{10, 20, 40, 80\}$ traders, with same $n/2$ number of buyers and sellers independently from the market mechanism and size. The computerized room where each classroom experiment was run had 90 trading positions (PCs), thereby allowing to simultaneously run at least two experimental sessions. More precisely, in each classroom: (i) only one of the two market mechanisms – DA or OTC – was implemented; (ii) market sizes were combined so that more than one market size was simultaneously implemented during the same experiment, with the total number of traders in each experiment being fixed and equal to the 90 trading positions (e.g., 1 market of size 80 together with 1 market of size 10, or

[1] Attanasi et al. (2016) used the same experimental procedures, softwares and comparable subject pools (cohorts from 2009 to 2014) to perform a first comparative study between OTC and DA markets with human traders and a fixed market size of 40 traders (for a different research question). Homogeneity across cohorts holds since Bocconi University undergraduate students' selection procedures and bachelor fields did not substantially change from 2009 to 2019.

1 market of size 40 with 2 markets of size 20 and 1 market of size 10, etc.). Table 1 shows the number of sessions (300 in total) implemented for each market mechanism-size combination, according to a between-subject design. The 8 market mechanism-size combinations in Table 1 also represents our 8 treatments.

	$n = 10$	$n = 20$	$n = 40$	$n = 80$
DA	80	40	20	10
OTC	80	40	20	10

Table 1: Number of Humans' experimental sessions for each market mechanism-size combination

Common features. The design of our classroom experiments follows the one in Attanasi et al. (2016), with the difference that we consider multiple market sizes. Here we summarize the design features that are treatment-independent:

- *Market and role assignment.* At the beginning of the experiment, each trader (sitting in front of one the 90 trading positions) is randomly assigned to one of the markets (sessions) set up for that experiment (e.g., to the 80-trader or to the 10-trader market). In each market, the traders are divided equally into buyers and sellers. Market and role assignment are kept constant during the whole experiment.

- *Number and length of trading periods.* Each *session* consists of three phases. Each *phase* consists of three trading periods of fixed length, where the length of a *period* is 120 seconds in the first two phases, and 60 seconds in the third phase.

- *Tradable units.* In each trading period, each seller owns one unit of a homogeneous good, and each buyer can purchase one unit of it. Therefore, as in Cason and Friedman (1996), subjects are allowed to trade only one unit per period. Hence, in each period of a market of size n, the maximum number of tradable units is $n/2$.

- *Redemption values.* At the beginning of each phase, each trader is assigned (exogenously as in Smith, 1962) a redemption value for the single unit he has to buy or sell. Buyers' redemption values are individual valuations of the good, and sellers' ones are their individual costs of production. More precisely, 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 receive for his/her unit.

- *Budget constraints.* A feasibility constraint is imposed in each period: buyers cannot bid over their own valuation, and sellers cannot ask under their own cost. Therefore, negative profits are not allowed. If a subject does not trade his/her commodity unit within the period, his/her profit is equal to zero.

- *Information.* At the beginning of the experiment, subjects learn the market size n and their role (either buyer or seller), which are both kept constant for the whole experiment. Then, at the beginning of each phase, subjects are given their redemption values, which are private information, are kept constant during the three periods of the phase and reshuffled at the end of it. Subjects do not know the distributions of valuations and costs in the market.[2] The last piece of information given to subjects is their ID number. While subjects' roles and redemption values remain fixed until the end of the experiment and the end of the phase, respectively, their ID is randomly reassigned at the beginning of every period. This prevents subjects from identifying trading counterparts in a given period on the basis of IDs learned in previous periods.

- *Incentives.* At the end of each trading period, each subject sees on the screen his/her profit as the difference between valuation and trading price – if he/she is a buyer – or between trading price and cost – if he/she is a seller. Being in a classroom experiment, subjects are not remunerated for their participation. However, they are given an incentive to play fairly: at the end of each 3-period phase, subjects are ranked according to their total profit in that phase.[3] Then, for each market of size n, we ask the $n/10$ traders having earned the highest and lowest total profit in that phase to stand up (e.g., the best 8 and worst 8 traders in markets with size 80, the best and the worst trader in markets with size 10). The former are praised publicly for their performance. The latter instead may be flouted by classmates.

Market mechanisms. As first treatment manipulation (see Table 1), we have two market mechanisms, DA and OTC:

- *DA markets.* Buyers and sellers post their bids and asks publicly, so that the bid-ask history of the market is public information: every agent is continuously informed about the highest bid and lowest ask standing in the market (*public trading*). The *bid/ask improvement rule holds*: to make a valid offer, an agent has to improve on the existing situation, i.e., a buyer (resp., seller) has to submit a bid higher (resp., ask lower) than the current highest bid (resp., lowest ask). When a buyer and a seller reach an agreement, they exit the market, the trading price appears on every subject's screen

[2]Note that, given the short amount of time of each trading period, we restricted the set of possible bids and asks to non-negative integer numbers lower than 100. Therefore, at the beginning of the experiment subjects are told that they can enter on the computer screen up to two-digit numers as bids and asks.

[3]Profits are in fact corrected to form the ranking: since redemption values are assigned randomly, subjects who are less lucky are penalized in terms of expected uncorrected profits. 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 phase profits will be corrected to elaborate the final ranking in the phase.

(*public information about trading prices*), the standing bids and asks are removed, and new bids and asks can be submitted.

- *OTC markets*. Each buyer (seller) can only send a bid (ask) 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 counterpart does not reply, the offer may be withdrawn and a new offer can be made that differs either in terms of the amount of the bid/ask, the counterpart's ID or both.[4] However, bids and asks submitted to a specific agent are not publicly disclosed (*private trading*), thus agents do not observe the highest bid and lowest ask standing in the market. The *bid/ask improvement rule does not hold*: the agents can always replace their current bid (ask) with a lower (higher) one. If an offer is privately accepted, the trading price appears on every subject's screen (*public information about trading prices*).

Market sizes. As second treatment manipulation (see Table 1), we implement markets of size n, with $n \in \{10, 20, 40, 80\}$, with $n/2$ buyers and $n/2$ sellers in each of these markets. 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. Independently of the market size n, the maximum buyers' valuation v and minimum sellers' cost c are set respectively at $\max v = 96$ and $\min c = 34$. The only difference between the four market sizes is in the distance between two subsequent valuations or costs, that is set at $80/n$. Therefore, there is a 8-integer, 4-integer, 2-integer, and 1-integer distance between two subsequent valuations or costs respectively in markets with 10, 20, 40, and 80 traders, with a finer grid of valuations and costs as n increases. This leads to the same equilibrium price ($p^* = 64$) independently of n, and a size-dependent efficient quantity $q^* = 0.8 \cdot n$, i.e., a size-independent ratio of efficient over total quantity: 80% of the available units are traded in equilibrium, independently of market size n.

2.2 Design of simulations with ZI agents

We implement agent-based simulations with two groups of Zero Intelligence (henceforth ZI) agents, buyers and sellers, by imposing the same market rules as in the experiments with human agents. To increase humans-automata comparability, we made these agent-based simulations by modelling ZI agents within the same z-Tree environment (Fischbacher, 2007) used to run the classroom experiments described in Section 2.1.

[4] If an agent 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.

Common features. All agents are homogeneous, with the only exception of the presence of heterogeneity in costs and valuations (for each market size n, same distribution as for human agents). Our ZI agents are Gode and Sunder (1993) zero intelligence traders with budget constraint: buyers can only bid in an interval between 0 and their own valuation; sellers can only ask in an interval between their cost and 100. Being random traders, bids and asks are drawn from a uniform distribution in the two respective intervals. The other main difference with respect to the experiments we have run with human agents is that – given that ZI agents have neither memory nor learning – we ran all trading periods from the first to the last one without organizing them in 3-period phases. Table 2 shows the number of periods that we ran for each of the 8 treatments. Note that, looking at Table 1, and considering 9 periods per session with human agents, the highest number of periods we ran with human agents is 720 (for both DA 10 and OTC 10), which is indeed lower than 800.

	$n = 10$	$n = 20$	$n = 40$	$n = 80$
DA	800	800	800	800
OTC	800	800	800	800

Table 2: Number of Automata's trading periods for each market mechanism-size combination

Market mechanisms. In the DA market, agents (randomly) choose only either their ask or bid; in the OTC market, agents (randomly) choose both their ask or bid and the counterpart to whom that ask or bid is going to be sent. In Sections 2.2.1 and 2.2.2 we report specific details of our DA and OTC random mechanisms.

Market sizes. We set the length of the trading period to be increasing with the market size. The length is 30 seconds for the markets of size 10, and it grows by a factor of 4 each time we double the market size, so that markets of size 20 have trading periods of 2 minutes length, those of size 40 have trading periods of 8 minutes length, and those of size 80 have trading periods of 32 minutes length. The choice of these parameters was due to the following theoretical reasons. First, we checked that 30 seconds was the minimum amount of time needed for DA markets of size 10 to reach the efficient quantity and obtain 0 intra-marginal inefficiency (as in Gode and Sunder, 1993). The DA market of size 10 is our baseline for ZI agents. Then we quadruplicate negotiation time every time the market size is doubled because the number of possible buyer-seller pairs grows by a factor of 4 when the market size doubles, and we take this as a proxy for the complexity of the interactions between agents. Keeping the time length of the trading period fixed across DA and OTC for the same market size provides a measure of OTC's intra-marginal inefficiency due to the greater complexity of OTC trading as compared to DA, i.e., only due to the market rules and not to (human) agents' strategic behavior.

2.2.1 Double Auction with ZI agents

We set a bid-ask improvement rule as in Brewer (2004, pp. 32-34). The buyer with the current best bid and the seller with the current best ask close the transaction. Then the market is reset to allow for more trades between the remaining ZI.

During each trading period, every 0.005 seconds *only* one ZI agent – randomly selected according to a uniform distribution – enters the market. The procedure that ZI agents follow in order to enter the market can be described according to subsequent rounds of 0.005 seconds each:

Round 1: The first ZI agent enters the market and generates a random ask, if it is a seller, or bid, if it is a buyer.[5] Then it becomes inactive and its offer stays on the market.

Round 2: Another ZI agent enters the market. Four cases can arise:

- The current offer on the market is a bid and the entering ZI agent is a buyer. It generates a new bid. If this new bid is higher than the current one, it becomes the new best bid; otherwise the agent exits and the market stays unchanged.

- The current offer on the market is a bid and the entering ZI agent is a seller. It generates a new ask. If the ask is lower or equal to the current bid, then the deal is closed at the current bid; the unit is traded, the two corresponding ZI agents are removed from the pool of possible traders and the market clears. If the ask is higher than the current bid, it remains as the best ask and the agent becomes inactive.

- The current offer on the market is an ask and the entering ZI agent is a seller. It generates a new ask. If this new ask is lower than the current one, it becomes the new best ask; otherwise the agent exits and the market stays unchanged.

- The current offer on the market is an ask and the entering ZI agent is a buyer. It generates a new bid. If the bid is higher or equal to the current ask, then the deal is closed at the current ask; the unit is traded, the two corresponding ZI agents are removed from the pool of possible traders and and the market clears. If the bid is lower than the current ask, it remains as the best bid and the agent becomes inactive.

Round 3: Another ZI agent enters the market. If the market has cleared in round 2, then round 3 is equivalent to round 1. If instead in the market an unmatched bid and/or ask are posted, two cases can arise:

[5] The random ask (or bid) generation process may involve multiple random draws. When this is the case, each new random draw happens after 0.005 seconds from the previous one.

- The new entering ZI agent is a buyer and generates a new random bid. If this bid is higher than the current bid and higher or equal to the current ask, the deal is closed at the current ask, the two corresponding ZI agents are removed from the pool of possible traders and the market clears. If the bid is higher than the current bid but lower than current ask, it simply replaces the best bid on the market. If the bid is lower than the current bid, nothing happens.

- The new entering ZI agent is a seller and generates a new random ask. If this ask is lower than the current ask and lower or equal to the current bid, the deal is closed at the current bid, the two corresponding ZI agents are removed from the pool of possible traders and the market clears. If the ask is lower than the current ask but higher than the current bid, it simply replaces the best ask on the market. If the ask is higher than the current ask, nothing happens.

Rounds $r > 3$: The market proceeds as in round 3 with rounds of 0.005 seconds each until there are no more available trades, or the trading period expires.

2.2.2 Over-the-Counter with ZI agents

We report here only the differences with respect to the DA treatment with ZI agents. The procedure that ZI agents follow in order to enter the OTC market can be described according to subsequent rounds (each round lasting 0.005 seconds, as for DA markets):

Round 1: All ZI agents send random offers. ZI buyers send a random bid to a ZI seller randomly selected according to a uniform distribution over the set of all ZI sellers; ZI sellers send a random ask to a ZI buyer randomly selected according to a uniform distribution over the set of all ZI buyers.

Round 2: Three cases can arise:

- ZI buyers that received one or more asks in round 1, compare their bid at round 1 with the lowest ask received. If the former is higher or equal to the latter, the unit is traded at a price equal to the lowest ask (as in Brewer 2004, Section 1.4.2), and the traders are removed from the market. Otherwise, the unit is not traded.

- ZI sellers that received one or more bids in round 1, compare their ask at round 1 with the highest bid received. If the former is lower or equal to the latter, the unit is traded at a price equal to the highest bid (as in Brewer 2004, Section 1.4.3), and the traders are removed from the market. Otherwise, the unit is not traded.

- ZI buyers and ZI sellers that did not receive any offers do nothing.

At the end of round 2, all bids and asks that did not end in a trade are cancelled. Thus, at the end of round 2, for ZI buyers and ZI sellers still on the market, the situation is identical to the situation at the beginning of round 1.

Odd Rounds $r \geq 3$: Though restricted to ZI buyers and ZI sellers still on the market, odd rounds $t \geq 3$ are identical to round 1, i.e., in odd rounds random offers are made to randomly selected counterparts.

Even Rounds $r \geq 4$: Though restricted to ZI buyers and ZI sellers still on the market, even rounds $t \geq 4$ are identical to round 2, i.e., in even rounds feasible trades are closed and, at the end of each even round, the remaining bids and asks are cancelled.

Rounds of 0.005 seconds each continue until there are no more ZI buyers or ZI sellers on the market, or the trading period expires.

3 Theoretical model and experimental hypotheses

This section aims at introducing the experimental hypotheses that will be tested in Section 4. As discussed in the Introduction, we can ground the experimental hypotheses on the relation between human agents and zero-intelligence agents on previous experimental studies on DA markets (see Gode and Sunder, 1993, 1997, 2004, 2018), and the ones on the effects of market size on efficiency of DA markets with humans on previous theoretical literature (see Chatterjee and Samuelson, 1983; Satterthwaite and Williams, 1989; Gresik and Satterthwaite, 1989; Wilson, 1985, 1987). However, there is no theoretical literature which allows us to formulate predictions on the effects of market size on the efficiency of OTC markets with human agents. For this reason, the section is divided in two parts. First, we put forward a simple and novel theoretical model of OTC markets which allows us to understand the strategic effects of a change in market size on human agents' behavior (Section 3.1). Then, we state the experimental hypotheses (Section 3.2).

3.1 A simple model of OTC markets

We propose a model of OTC market in which sellers and buyers meet to trade a good. In the literature, OTC markets have been modeled by Duffie et al. (2007) and follow-up papers.[6] The model by Duffie et al. (2007) is a search model where agents with high and low valuations of an asset want to trade it. However, it has a limitation that becomes particularly relevant when comparing the efficiency of OTC markets with human agents to the efficiency of OTC markets with ZI agents with varying market size: in their model,

[6] Duffie (2010, 2012); Ashcraft and Duffie (2007); Duffie et al. (2005, 2007); Duffie and Manso (2007); Duffie et al. (2009, 2010a,b, 2014).

whenever a profitable match occurs, there is a trade. As a consequence, the search model does not take into account the fact that individual bargaining strategies in every period are affected by expectations of future payoffs, which in turn are affected by the market size. We therefore introduce a bargaining model of OTC markets that is simpler than the search model by Duffie et al. (2007) but that explicitly models this feature of the strategic interaction, which mostly distinguishes human agents from ZI agents.

Specifically, we model the bargaining process as a sequence of take-it-or-leave-it offers. Each seller has one unit to sell, and each buyer is willing to purchase one unit only. We consider the market size $n = 10$, with 5 buyers and 5 sellers. Buyers privately evaluate the object $v \in \{0, 1, 2, 3, 4\}$, while sellers have a cost of production $c \in \{0, 1, 2, 3, 4\}$. Buyers' valuations (and sellers' costs) are their private information. As in our experiment, there is exactly one buyer (resp., seller) with each valuation (resp., cost), so that each valuation (resp., cost) is equally likely.[7] Let $\beta \in \{0, 1, 2, 3, 4\}$ be the type of buyer, and $\sigma \in \{0, 1, 2, 3, 4\}$ be the type of seller. With a slight abuse of notation, we call β (resp., σ) not only buyers' (resp., sellers') types, but also their valuations (resp., costs).

We consider the following sequential bargaining protocol. Agents interact for 2 periods.[8] All the traders are present in the market at the beginning of the game. In every period each buyer is randomly matched to a seller (the market is symmetric, same number of buyers and sellers). Given the match, each trader is selected with probability 1/2 to be the proposer. The selection of the proposer happens independently in every match. The proposer makes a take-it-or-leave-it offer. If the offer is accepted the good is traded and the two traders leave the market. If the offer is rejected the partnership dissolves and the traders move to the following period in which a new random matching occurs. Proposition 1 characterizes a subgame perfect equilibrium of the game. The proof is contained in Appendix A.

Proposition 1 *The following is a subgame perfect equilibrium of the two-period bargaining game.*

- *Period* 1 *offers are:*

$$p_1^\beta = \begin{cases} 0 & \text{if } \beta \in \{0,1\} \\ 1 & \text{if } \beta = 2 \\ 2 & \text{if } \beta \in \{3,4\}, \end{cases} \qquad p_1^\sigma = \begin{cases} 2 & \text{if } \sigma \in \{0,1\} \\ 3 & \text{if } \sigma = 2 \\ 4 & \text{if } \sigma \in \{3,4\}. \end{cases}$$

[7] The theoretical analysis is the same even if we assume that buyers' valuation and sellers' cost are randomly drawn from a discrete uniform over $\{0, 1, 2, 3, 4\}$. We chose to model the market as described above, because it corresponds more closely to the specifics of our experimental environment.

[8] Note that what we call period in the model is not the trading period of the experiment. The trading period in the model is the union of the two periods in which the agents interact.

- *Period* 1 *acceptance strategies are:*

$$s_2^\beta = \begin{cases} \textit{Yes} & \textit{if } p_2 \leq \max\{0, \beta - 1\} \\ \textit{No} & \textit{otherwise;} \end{cases} \qquad s_2^\sigma = \begin{cases} \textit{Yes} & \textit{if } p_2 \geq \min\{\sigma + 1, 4\} \\ \textit{No} & \textit{otherwise.} \end{cases}$$

- *Period* 2 *offers are:*

$$p_2^\beta = \begin{cases} 0 & \textit{if } \beta \in \{0,1\} \\ 1 & \textit{if } \beta \in \{2,3\} \\ 2 & \textit{if } \beta = 4, \end{cases} \quad p_2^\sigma = \begin{cases} 2 & \textit{if } \sigma = 0 \\ 3 & \textit{if } \sigma \in \{1,2\} \\ 4 & \textit{if } \sigma \in \{3,4\}. \end{cases}$$

- *Period* 2 *acceptance strategies are:*

$$s_2^\beta = \begin{cases} \textit{Yes} & \textit{if } p_2 \leq \beta \\ \textit{No} & \textit{otherwise;} \end{cases} \qquad s_2^\sigma = \begin{cases} \textit{Yes} & \textit{if } p_2 \geq \sigma \\ \textit{No} & \textit{otherwise.} \end{cases}$$

Note that in period 1 buyers and sellers strategically reject more offers than in period 2, because they anticipate they may have higher payoffs in the future, either because they will be the proposers, or because they will end up in a better match. This strategic waiting effect decreases the number of trades in the first period. Furthermore, offers are also higher in period 1 than in period 2, to compensate the effect of lower acceptance rates. As a matter of fact, trades happen in the cases described in Table 3.

Period	*Proposer*	(β, σ) pairs who trade
1	Buyer	$\{(2,0),(3,0),(4,0),(3,1),(4,1)\}$
	Seller	$\{(3,0),(4,0),(3,1),(4,1),(4,2)\}$
2	Buyer	$\{(0,0),(1,0),(2,0),(3,0),(4,0),(2,1),(3,1),(4,1),(4,2)\}$
	Seller	$\{(2,0),(3,0),(4,0),(3,1),(4,1),(3,2),(4,2),(4,3),(4,4)\}$

Table 3: Feasible trades, by periods and role of the proposer

Effects of market size. Our model suggests a channel through which market size may affect efficiency and volumes of trade in OTC markets. Recall that, as discussed in the Introduction, Wilson (1987) tackles the theoretical analysis of the efficiency of DA markets as their size grows, showing that they converge to efficiency. Wilson (1987) argues that this is due to an increase in the competitive pressure by the same type of traders which is particularly effective in DA markets, where offers are made publicly to all traders of one side of the market. In OTC markets, this effect does not apply. As a matter of fact, offers in OTC markets are made only to a single trader. Therefore, the trader who receives the offer does not compete with other traders of his/her type (buyers or sellers) for that specific offer. Market size therefore has the opposite effect. Instead of pushing traders to accept advantageous offers more quickly, it increases their expected gains from future trades, and induces them to wait longer and/or for better offers. Therefore, only more advantageous trades happen in earlier periods, and this reduces the efficiency of the market.

To see this we can compare the $n = 10$ market described above with an $n = 6$ market where we have three buyers with valuations $\beta \in \{0, 2, 4\}$ and three sellers with costs $\sigma \in \{0, 2, 4\}$. Appendix B contains the formal analysis of both the $n = 6$ market and the comparison between markets of size 6 and 10. This comparison of the two markets highlights the two mechanisms at play: first, as market size increases traders are less willing to accept low offers in the first period; second, traders anticipate this effect and increase first-period offers when the market size increases. Overall, the first effect dominates, and we find that, when moving from the smaller to the larger market, the probability of observing a number of trades lower than the efficient one increases, and expected efficiency, measured as total surplus over efficient surplus, decreases. This analysis informs our experimental predictions of Section 3.2.

3.2 Experimental hypotheses

We structure our experimentally testable predictions in the following way: first, we consider the comparison of the performance of OTC and DA markets by type of traders and market size in terms of traded quantity (**H1**), trading price (**H2**), efficiency (**H3**), and sources of inefficiency (**H4**). Then we follow the theoretical considerations of the model of Section 3.1 to formulate an experimental hypothesis on the effects of market size on OTC markets with humans (**H5**).

Traded quantity. The first market outcome that we consider is traded quantity. For human traders, we expect that, consistently with previous evidence (Attanasi et al., 2016), OTC markets with humans induce smaller volumes of trade than DA markets, for each given market size. Our theoretical model suggests that bigger OTC markets with humans should more often induce inefficient volumes of trade, specifically they have a *higher* likelihood of observing a traded quantity lower than the efficient one, due to the strategic effects on acceptance rates of earlier offers. On the contrary, the experimental literature suggests that DA markets with humans tend to efficiency when market size increases (Friedman and Rust, 1993; Friedman, 2018). Therefore, with DA we expect a *lower* likelihood of observing a traded quantity lower than the efficient one as market size increases.

For automata, results from agent-based simulations under DA (for a review, see Gode and Sunder, 2018; Rust et al., 2018) show traded quantities higher than the efficient one more often as market size increases. We hypothesize that the same holds under OTC, due to our agent-based simulations with length of trading periods increasing with the market size so as to compensate for increased complexity of interactions among ZI agents. However, due to the greater complexity of the OTC mechanism, we expect lower traded quantity than under the DA mechanism, independently from the market size. Hypothesis **H1** summarizes our predictions on traded quantity.

H1: (i) [*Automata*] OTC markets have lower traded quantity than DA markets for each market size, and traded quantity of both markets is more often higher than the efficient one as market size increases.

(ii) [*Humans*] OTC markets have lower traded quantity than DA markets for each market size, and traded quantity of OTC (resp., DA) markets is more (resp., less) often lower than the efficient one as market size increases.

(iii) [*Humans vs. Automata*] For OTC markets, the humans-automata gap in terms of traded quantity is increasing in market size. This does not occur in DA markets.

Trading price. Following the results of Attanasi et al. (2016) for markets of size 40, we expect OTC markets with humans to have lower trading prices than DA markets, and higher standard deviation around the equilibrium price. Moreover, we expect trading prices to converge to the equilibrium price as market size increases in DA markets (Friedman and Rust, 1993; Friedman, 2018), but not in OTC markets, consistently with the theoretical discussion of Section 3.1. For automata, we know that ZI agents never converge to the equilibrium price, independently from the market size. Therefore, we expect humans to perform increasingly better – lower distance between actual and equilibrium price – than automata under DA as market size increases. According to our theoretical predictions for OTC markets, the opposite comparative statics should hold, with humans possibly performing worse than automata for large enough market sizes. Hypothesis **H2** summarizes our predictions on trading prices.

H2: (i) [*Automata*] The distance between trading price and equilibrium price is independent from the market trading mechanism (DA or OTC) and the market size.

(ii) [*Humans*] OTC markets have lower trading prices than DA markets, trading prices of DA (resp., OTC) markets converge (resp., do not converge) to the equilibrium price as market size increases.

(iii) [*Humans vs. Automata*] For OTC (resp., DA) markets, the humans-automata difference in the distance between trading price and equilibrium price increases (resp., decreases) with market size.

Efficiency. The experimental predictions on efficiency are grounded on different sources. For what concerns markets with automata, we rely on the same experimental literature for DA and the same design features for OTC that we discussed when we elaborated H1 above. In particular, results in the literature on agent-based simulations under DA show increasingly higher efficiency when the market size increases, and our OTC markets with ZI are endowed with the same trading environment to reach the same positive efficiency - size correlation. However, given greater complexity, they should be less efficient than DA ones. For what concerns markets with humans the existing theoretical literature on DA markets shows that efficiency should be increasing with market size; our theoretical model suggests that efficiency of OTC markets is decreasing with market size, and we expect OTC markets to be less efficient than DA markets, as in Attanasi et al. (2016). All these predictions are summarized in hypothesis **H3**.

H3: (i) [*Automata*] OTC markets are less efficient than DA markets, and efficiency of both types of markets is increasing with market size.

(ii) [*Humans*] Efficiency of DA (resp., OTC) markets is increasing (resp., decreasing) with market size, and the DA-OTC efficiency gap is positive and increasing with market size.

(iii) [*Humans vs. Automata*] For each market size, the DA-OTC efficiency gap is lower for ZI than for human agents.

Sources of inefficiency. Following Gode and Sunder (1993) and Attanasi et al. (2016), we decompose the loss of efficiency into two possible sources of inefficiency. We distinguish between the inefficiency that comes from extra-marginal units being traded (*EM-inefficiency*) and the inefficiency that comes from intra-marginal units not being traded (*IM-inefficiency*). As for automata, Gode and Sunder (1993) run agent-based simulations of DA markets that allow all ZI intra-marginal traders to trade, i.e., with no IM-inefficiency. We did the same for our DA markets, independently of the size, and maintained the same length of the trading period for the OTC market of the correspondent size, which may lead to some IM-inefficiency independently from the market size. This effect is due to a greater trading complexity with respect to DA, as trading complexity is higher under OTC for each market size. As for human subjects, Attanasi et al. (2016) show that inefficiency of DA markets is mostly associated with EM-inefficiency, while the inefficiency of OTC markets is a mixture of the two. We investigate the effects of market size on the composition of market inefficiency. We expect IM-inefficiency to disappear from DA markets as their size increases, while we do not expect the same to happen for OTC markets, due to their one-to-one trading mechanism. Hypothesis **H4** summarizes this prediction.

H4: (i) [*Automata*] DA markets have no IM-inefficiency, independently from the market size. OTC markets have IM-inefficiency and its share over EM-inefficiency is independent from the market size.

(ii) [*Humans*] The share of IM-inefficiency over EM-inefficiency is decreasing with market size in DA markets, while this does not occur in OTC markets.

(iii) [*Humans vs. Automata*] For DA markets, the humans-automata gap in terms of positive IM-inefficiency is decreasing in market size. This does not occur in OTC markets.

Strategic effects on OTC markets with humans. As discussed in the model of OTC markets with humans in Section 3.1, the decrease in efficiency of OTC markets as market size increases – postulated in H3(ii) – should be due to a strategic effect of OTC markets. Contrarily to DA markets (Wilson, 1987), in OTC markets an increase in size is not associated to higher competitive pressure from same-type traders, while it is associated, at least in earlier periods, to higher expected continuation profits if a trader does not accept an offer. This means that acceptance rates of traders who receive an offer are lower in earlier periods, as market size grows. This can be shown by testing hypothesis **H5**.

H5: [*Humans*] In OTC markets only the most advantageous offers are accepted in the earlier periods, and this effect is stronger as market size increases. All this does not hold in DA markets.

4 Results

The analysis of the results follows the order of the experimental hypotheses (Section 3.2). First, we compare the performance of DA and OTC markets by type of traders and market size in terms of traded quantity (Result 1), trading price (Result 2), efficiency (Result 3), and sources of inefficiency (Result 4). We then focus on OTC markets. We notice that OTC markets with humans respond to market size differently from OTC markets with ZI agents and we investigate whether this phenomenon is caused by the strategic mechanisms analyzed theoretically in Section 3 (Result 5).

Result 1: Traded quantity. Table 4 shows the percentage of periods in which the traded quantity is lower, equal or higher than the efficient one, by type of traders, treatment and market size. Table 5 reports the results of probit regressions where the dependent variables indicate whether the traded quantity is, respectively, below (odd columns) and above (even columns) the efficient one. Data from automata and humans are analyzed separately (left panel and right panel of the table, respectively). The explanatory variables are: a treatment dummy for OTC markets (with DA as baseline), three dummies for market size (with size 10 as baseline), and the interaction between the treatment dummy and the dummies for market size.

	Automata			**Humans**		
	$q < q^*$	$q = q^*$	$q > q^*$	$q < q^*$	$q = q^*$	$q > q^*$
DA						
10 agents	0	31.9	68.1	33.7	49.6	16.7
20 agents	0	11.5	88.5	29.4	47.8	22.8
40 agents	0	3.5	96.5	20.0	28.9	51.1
80 agents	0	1.5	98.5	13.3	11.1	75.6
OTC						
10 agents	0.3	43.2	56.5	56.4	32.4	11.2
20 agents	2.3	34.1	63.6	48.3	30.0	21.7
40 agents	7.1	23.9	69.0	57.2	23.9	18.9
80 agents	5.9	14.7	79.4	86.7	7.8	5.5

Table 4: Percentage of periods in which the traded quantity is lower, equal or higher than the efficient one, by type of traders, treatment and market size.

	Automata		**Humans**	
	$q < q^*$	$q > q^*$	$q < q^*$	$q > q^*$
OTC	3.276	-0.307***	0.580***	-0.246***
	(144.51)	(0.06)	(0.07)	(0.08)
20	0.000	0.729***	-0.121	0.221**
	(204.37)	(0.07)	(0.08)	(0.09)
40	0.000	1.340***	-0.422***	0.995***
	(204.37)	(0.10)	(0.12)	(0.11)
80	0.000	1.699***	-0.691***	1.659***
	(204.37)	(0.12)	(0.17)	(0.15)
OTC × 20	0.802	-0.544	-0.081	0.208
	(204.37)	(0.10)	(0.11)	(0.13)
OTC × 40	1.340	-1.008***	0.443***	-0.663***
	(204.37)	(0.11)	(0.15)	(0.16)
OTC × 80	1.242	-1.043***	1.641***	-2.039***
	(204.37)	(0.14)	(0.24)	(0.27)
Const.	6.083	-0.471***	-0.419***	-0.967***
	(144.51)	(0.05)	(0.05)	(0.05)

Table 5: Probit regression of a dummy for quantity below and above the efficient one over treatment and market size variables, separately for ZI agents and human agents (standard errors in brackets); * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

With **Automata**, in DA markets the traded quantity is never lower than the efficient one.[9]

[9]This is a combination of efficiency of DA markets with automata and the fact that we chose the length

Table 4 also shows that the traded quantity is usually higher than the efficient quantity in OTC markets (χ^2 test for $q > q^*$ vs. $q \leq q^*$, *p-value* < 0.001 for each market size), although it is always significantly lower than in DA markets independently of the size. The latter is confirmed by the significance of the negative coefficient of the OTC dummy for $q > q^*$ in Table 5. We also notice that the negative difference in the traded quantity between OTC and DA increases with market size, since the negative coefficient of dummy OTC $\times$ 20 for $q > q^*$ is not significant, while those of dummies OTC $\times$ 40 and OTC $\times$ 80 for $q > q^*$ are. However, as market size increases, the percentage of periods with traded quantity greater than the efficient quantity significantly increases also for OTC markets (+23 percentage points moving from size 10 to size 80 in Table 4, increase significant at the 0.1% level, χ^2 test). In confirmation of that, for $q > q^*$, the coefficients of dummies 20, 40 and 80 are positive, significant and significantly higher, in absolute value, than those of dummies OTC $\times$ 20, OTC $\times$ 40 and OTC $\times$ 80, respectively (Wald test of difference of coefficients, *p-value* < 0.01 in the three cases). With this, we can conclude that *H1(i) is confirmed.*

With **Humans**, contrarily to automata, DA and OTC markets respond differently to market size – see the right panel of Table 4, especially column $q > q^*$ for DA (+30 percentage points moving from size 10 to size 80, increase significant at the 0.1% level, χ^2 test) and column $q < q^*$ for OTC (+49 points moving from size 10 to size 80, increase significant at the 0.1% level, χ^2 test). The third and fourth column of Table 5 confirm that the probability of being below (resp., above) the efficient quantity increases (resp., decreases) in OTC markets, since the coefficient of the OTC dummy is positive for $q < q^*$ and negative for $q > q^*$, significant in both cases. Moreover, in DA markets the probability of observing traded quantity lower than the efficient one decreases with market size, significantly for market size $n \in \{40, 80\}$. OTC markets display the opposite behavior, which is captured by the treatment dummy interacted with the dummies for market size, again significant for market size $n \in \{40, 80\}$. With this, we can conclude that *H1(ii) is confirmed.*

As for the **comparison between Humans and Automata**, Table 4 shows that under OTC the humans-automata (resp., automata-humans) difference in the probability of observing $q < q^*$ (resp., $q > q^*$) increases monotonically and significantly from 56 (resp., 42) percentage points for markets of size 10 to 81 (resp., 74) percentage points for markets of size 80 (t-test of the difference in percentages, *p-value* < 0.01). Therefore, for OTC markets, the humans-automata gap in terms of traded quantity is increasing in market size. As for DA markets, Table 4 shows the opposite picture, with the humans-automata (resp., automata-humans) difference in the probability of observing $q < q^*$ (resp., $q > q^*$) decreasing monotonically

of the trading period for DA markets of size 10 with the specific aim of allowing enough time to automata to reach the efficient quantity (see Gode and Sunder, 1993). In this way, the DA market of size 10 works as a baseline for the performance of DA markets of larger size and OTC markets of any size (see subsection "Market sizes" in the design of simulations with ZI agents of Section 2.2).

from 34 (resp., 51) percentage points for markets of size 10 to 13 (resp., 23) percentage points for markets of size 80 (t-test of the difference in percentages, *p-value* < 0.01). Therefore, for DA markets, the humans-automata gap in terms of traded quantity is decreasing in market size. With this, we can conclude that *H1(iii) is confirmed.*

Result 2: Trading prices. We now turn to the analysis of the difference between trading prices and the equilibrium price. Table 6 reports the average distance of the trading prices from the equilibrium price, by type of traders, treatment and market size.

	Automata		**Humans**	
	DA	**OTC**	**DA**	**OTC**
10 agents	0.35	0.65	1.71	−0.44
	(0.23)	(0.22)	(0.14)	(0.14)
20 agents	1.28	0.89	0.14	−1.59
	(0.16)	(0.15)	(0.11)	(0.12)
40 agents	1.72	1.08	0.42	−2.55
	(0.12)	(0.11)	(0.11)	(0.12)
80 agents	1.58	0.70	−0.03	−3.31
	(0.09)	(0.08)	(0.12)	(0.13)

Table 6: Average distance of trading prices from the equilibrium price (standard errors in brackets).

With **Automata**, the average trading price is significantly higher than the equilibrium price for almost all treatment-size combinations, with the exception of DA markets with 10 agents where it is higher but not significantly so (t-test, *p-value* $= 0.128$). Furthermore, the average trading price is not significantly different between DA and OTC markets for size $n = 10$ (t-test, *p-value* $= 0.35$), while it is significant at the 10% level for $n = 20$ and at the 1% level for $n \geq 40$. More precisely, we detect a higher DA-OTC spread in the average trading price as market size increases. Finally, the average trading price is not invariant to the market size for DA markets (t-test, *p-value* $= 0.002$), while it is invariant to the market size for OTC markets (t-test, *p-value* $= 0.106$). With this, we can conclude that *H2(i) is only partially confirmed.*

With **Humans**, for DA markets, the average trading price is significantly higher than the equilibrium price for market size 10 (t-test, *p-value* < 0.001). As market size increases, the positive difference between the average trading price and the equilibrium price decreases and the former basically coincides with the latter for the highest market size (t-test, *p-value* $= 0.773$). This is consistent with the experimental literature on DA markets which suggests that DA markets converge to competitive equilibrium as market size increases (see Chatterjee and Samuelson (1983) and follow up papers). On the contrary, we observe that the average trading price is significantly lower than the equilibrium price in OTC markets independently of the size n (*p-value* < 0.001 for each n), and that this difference is increasing in market size (*p-value* < 0.01 for each comparison 10 vs. 20, 20 vs. 40, and 40 vs. 80). The latter is further evidence that in OTC markets human agents perform increasingly worse as market size grows. With this, we can conclude that *H2(ii) is confirmed.*

In the **comparison between Humans and Automata**, for OTC markets the humans-automata difference in the average distance between trading price and equilibrium price increases monotonically from 0.11 for markets of size 10 to 2.61 for markets of size 80 (t-test for the difference in the trading-equilibrium price distance between humans and automata in markets with 80 agents vs. markets with 10 agents is significant at the 1% level, *p-value* < 0.01). As for DA markets, we find an opposite trend, with the humans-automata difference in the average distance between trading price and equilibrium price decreasing monotonically and significantly from 1.36 for markets of size 10 to -1.55 for markets of size 80 (t-test for the difference in the trading-equilibrium price distance between humans and automata in markets with 80 agents vs. markets with 10 agents is significant at the 1% level, *p-value* < 0.01). Overall, the analysis of the comparison between humans and automata let us to state that *H2(iii) is confirmed.*

Result 3: Efficiency. Let us now focus on market efficiency itself. We define the Market Efficiency Index as the surplus realized from trade over the potential surplus. Table 7 reports the values of the efficiency index by type of traders, treatment and market size. Table 8 reports results of a beta regression of the efficiency index by type of traders, treatment and market size.

	Automata		**Humans**	
	DA	**OTC**	**DA**	**OTC**
10 agents	99.08%	98.74%	93.2%	85.6%
20 agents	97.85%	97.57%	93.2%	87.2%
40 agents	97.60%	97.16%	94.9%	87.0%
80 agents	97.62%	97.07%	94.7%	82.6%

Table 7: Efficiency index, by type of traders, treatment and market size.

With **Automata**, efficiency is higher under DA than under OTC for each market size (Mann-Whitney test; only non-significant difference detected for market size 10, *p-value* $= 0.108$, second highest *p-value* for market size 20, *p-value*$= 0.001$). By looking at the first two columns of Table 8, we notice that efficiency in DA markets decreases in deviations from the efficient quantity q^*: this is a direct effect of what observed in Result 1, i.e., that DA markets with ZI agents have often higher than equilibrium trading volumes (see Table 4). However, in contrast with our predictions, Table 7 shows that market efficiency, both under DA and under OTC, is significantly lower in markets of size $n \geq 20$ than in the smallest market of size $n = 10$, and essentially constant across market sizes $n \in \{20, 40, 80\}$ (this is confirmed

	Automata		Humans	
	DA	**OTC**	**DA**	**OTC**
$(q - q^*)/q^*$	-8.559	0.360	3.658	3.676
	(0.000)	(0.007)	(0.000)	(0.000)
$\lvert\bar{p} - p^*\rvert$	0.002	-0.009	0.009	-0.003
	(0.549)	(0.086)	(0.209)	(0.593)
20	-1.160	-0.240	0.310	-0.105
	(0.000)	(0.000)	(0.000)	(0.013)
40	-1.621	-0.332	0.438	-0.127
	(0.000)	(0.000)	(0.000)	(0.017)
80	-1.774	-0.417	0.244	-0.297
	(0.000)	(0.000)	(0.003)	(0.000)
Seller Sender	0.003	0.028	0.044	-0.165
	(0.952)	(0.667)	(0.604)	(0.027)
Const.	6.338	3.852	2.323	2.529
	(0.000)	(0.000)	(0.000)	(0.000)

Table 8: Beta regressions of efficiency by treatment over all trading periods (*p-values* in brackets).

by non-significant differences among coefficients of dummy variables 20, 40 and 80 in the regressions of Table 8, both under DA and under OTC). Importantly, as noticed for Results 1 and 2, in the case of ZI agents DA and OTC markets exhibit the same comparative statics in terms of market size. Overall, we can conclude that *H3(i) is only partially confirmed.*

With **Humans**, we first confirm that DA markets are more efficient than OTC markets for each market size (Mann-Whitney test; *p-value* < 0.001 independently of the market size), consistently with previous experimental evidence (see Attanasi et al. 2016). Moreover, by looking at the third and fourth columns of Table 8, we confirm that DA markets and OTC markets behave in a very different fashion also in terms of efficiency. In fact, while the dummy variables of market sizes 20, 40 and 80 have a significant positive effect on the DA market efficiency, the opposite holds for OTC market efficiency. Furthermore, in OTC market the fact that the sender of the accepted offer is a seller rather than a buyer decreases market efficiency (significant and negative coefficient of dummy variable Seller Sender in the fourth column of Table 8). In fact, this is a signal of sellers' pressure that makes them increase their willingness to sell, ultimately leading to lower than equilibrium trading prices reported in Result 2 for each n-size OTC market with humans (see the last column of Table 6). The latter dummy variable is not significant in the DA treatment. Therefore, we can conclude that *H3(ii) is confirmed.*

As for the **comparison between Humans and Automata**, Table 7 shows that the DA-OTC efficiency gap is lower for automata (first two columns) than for human agents (last two columns). The hypothesis that the DA-OTC efficiency gap is lower for automata than for human agents cannot be rejected for any market size (t-test, lowest *p-value*$= 0.999$).

Therefore, we conclude that *H3(iii) is confirmed.*

Result 4: Sources of Inefficiency. We now focus on the two possible sources of inefficiency. As anticipated in Section 3.2, we distinguish between *EM-inefficiency* (extra-marginal units being traded) and *IM-inefficiency* (intra-marginal units not being traded). Table 9 shows the composition of inefficiency by treatment, type of traders and market size.

	Automata		**Humans**	
	EM-Ineff.	IM-Ineff.	EM-Ineff.	IM-Ineff.
DA				
10 agents	100.0	0.0	71.8	28.2
20 agents	100.0	0.0	74.8	25.2
40 agents	100.0	0.0	90.6	9.4
80 agents	100.0	0.0	94.9	5.1
OTC				
10 agents	99.0	1.0	51.1	48.9
20 agents	98.0	2.0	65.0	35.0
40 agents	96.3	3.7	67.6	32.4
80 agents	98.0	2.0	49.1	50.9

Table 9: Share of sources of inefficiency: extra-marginal and intra-marginal.

With **Automata**, IM-inefficiency plays no role under DA for each market size. This is by construction for DA markets of size 10, for which we chose the length of the trading period so as to allow enough time to automata to reach the efficient quantity (see Gode and Sunder, 1993). Absence of IM-inefficiency also in DA markets with $n \geq 20$ confirms that the exogenous increase of the negotiation time with market size, which was meant to account for the increase in the number of possible buyer-seller pairs, was able to compensate for the whole increase in complexity of agents' interactions under the DA trading mechanism. We detect instead a positive though small IM-inefficiency for OTC markets, which is however significantly different from zero for each market size (t-test for the mean IM-inefficiency equal to 0 for each market size: null hypothesis is rejected at 5% level for every market size; *p-value* $= 0.014$ for market size 10; *p-values* < 0.01 for any other market size). This is proof of a slightly greater complexity of the OTC trading mechanism as compared to the DA one, given the same market size and the same allowed trading time. However, as market size increases, the weight of IM-inefficiency over EM-Inefficiency is monotone in size up to $n = 40$. In fact, t-tests of the difference in IM-inefficiency for increasing market size reveal a significant positive difference for 20 vs. 10 and 40 vs. 20 (*p-values* < 0.01). However, IM-inefficiency significantly decreases between 40 and 80 (*p-value* < 0.01). Therefore, we can conclude that *H4(i) is only partially confirmed.*

With **Humans**, the weight of IM-inefficiency over EM-inefficiency decreases dramatically with the number of agents in DA. In fact, differences in the share of IM-inefficiency between a market size and a smaller one are all negative for DA markets. We use a difference-in-means test to assess whether these differences are statistically significant. At the 1% level of significance, the differences are not equal to zero in DA markets, with the exception of the difference between the market of size 20 and the one of size 10, which is negative but not significant (*p-value*$=0.519$). As for OTC markets, the weight of IM-inefficiency over EM-inefficiency does not increase monotonically with market size. Furthermore, a difference-in-means test assesses that the difference between size 80 and size 10 is positive and significant at the 1% level (*p-value*< 0.01), and that the difference between size 80 and size 40 is positive and significant at the 1% level (*p-value*< 0.01). Therefore, for OTC markets, IM-inefficiency does not disappear at market size increases. This is once again evidence of the strategic differences of the two trading mechanisms when used by humans. With this, we conclude that *H4(ii) is confirmed.*

As for the **comparison between Humans and Automata**, for DA markets the humans-automata gap in terms of IM-inefficiency is positive for each market size, and it decreases significantly with market size. This directly follows from the fact that IM-inefficiency is null for automata for each market size (second column of Table 9) and that it is positive and significantly decreasing in the market size for humans (fourth column of Table 9). As for OTC markets, the humans-automata gap in terms of IM-inefficiency is positive but not decreasing with market size. Again, this directly follows from the fact that IM-inefficiency is small and constant for automata for each market size (second column of Table 9) and that it is positive but quite stable in the market size for humans (fourth column of Table 9), even significantly higher for market size 80 than for any other market size (highest *p-value*$= 0.0003$ for the comparison with market size 10). With this, we conclude that *H4(iii) is confirmed.*

Result 5: Strategic effects in OTC with humans. Let us now focus on the strategic features that distinguish OTC markets from DA markets. In Section 3 we argued that OTC markets with human agents may be exposed to an increase in inefficiency due to the strategic behavior of traders who are less willing to accept offers in earlier periods because they anticipate possibly higher gains from trade in future periods. We summarized this intuition in H5, which states that acceptance rates of traders who receive an offer are lower in earlier periods, as market size grows. Table 10 shows the difference between trading prices in the first $n/5$ trades in a period, by treatment and market size n, with $n \in \{10, 20, 40, 80\}$. We note that, in OTC markets, this difference is decreasing in market size (second column of Table 10, *p-value* < 0.01 for each comparison 10 vs. 20, 20 vs. 40, and 40 vs. 80). This information, together with the observation that the majority of the offers leading to a transaction are proposed by sellers (between 55% and 61% depending on market size, significantly higher than 50% at the 1% level) supports our claim that under OTC only the most advantageous offers are accepted in the earlier trading periods, and that this effect is stronger as market size increases.

	DA	**OTC**
10 agents	0.81	-0.93
	(0.28)	(0.28)
20 agents	-0.42	-2.26
	(0.17)	(0.16)
40 agents	0.02	-3.22
	(0.18)	(0.17)
80 agents	-0.04	-4.02
	(0.20)	(0.18)

Table 10: Difference between trading prices for the first $n/5$ transactions, where $n \in \{10, 20, 40, 80\}$ is the number of agents in the market (standard errors in brackets).

As a control, the first column of Table 10 shows that the same does not occur for DA markets. In fact, the difference between trading prices for the first $n/5$ transactions does not decrease in market size: it is significantly higher than 0 for market size 10 (*p-value*< 0.01), significantly lower than 0 for market size 20 (*p-value*< 0.01), and basically null for the two biggest market sizes (lowest *p-value*$= 0.849$). Furthermore, the fraction of offers leading to a transaction due to sellers' asks is not significantly higher than the fraction due to buyers' bids (t-test of the difference of the average fraction of transactions is equal to 0, *p-value*$= 0.447$). Hence, we conclude that *H5 is confirmed.*

5 Conclusions

In this paper we analyze the behavior of over-the-counter (OTC) markets of varying size both when the interaction occurs among humans and when the interaction occurs among zero-intelligent (ZI) agents. To the best of our knowledge, this is the first theory-driven experimental study of OTC markets that compares such decentralized market mechanism with centralized trading institutions (DA), controlling for human vs. automata performance under different market sizes.

We explore electronic OTC mechanisms with public information about trading prices. We think that nowadays such trading institutions have significant economic applications. In a large number of financial markets, negotiations and transactions occur on a bilateral basis rather than, as in auction markets, through publicly posted bids and asks. Moreover, these negotiations and transactions occur via computer rather than, as in pit markets, orally. Many types of government and corporate bonds, real estate, currencies, and bulk commodities are typically traded electronically over the counter. Furthermore, in a number of these 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 discussion of the economic relevance of OTC markets, see, e.g., Duffie et al., 2005, Duffie, 2012).

The first contribution of our study is theoretical: in the literature, OTC markets have been mainly framed as search models with agents with high and low valuations who want to trade an asset (see Duffie et al., 2007, and follow-up papers). We think that this approach disregards some key strategic features of OTC markets. To account for these features, we introduce a simple bargaining model which explicitly describes agents' strategic interaction in OTC markets: strategic interaction mostly distinguishes human agents from ZI agents' behavior. In fact, humans' bargaining strategies in every period are affected by expectations of future payoffs, which in turn are affected by market size. Therefore, our model is especially apt to analyze how OTC markets' trading and efficiency evolve as market size increases.

Relying on a large experimental dataset of 6,400 undergraduate students in Economics and Management, we show that the main predictions of our model – lower trading quantity, higher price dispersion and lower market efficiency when OTC market size increases – are essentially verified. Furthermore, with human traders the DA-OTC efficiency gap increases as market size increases. Finally, agent-based simulations with ZI agents show that these results are not due to OTC market rules. In fact, OTC markets with ZI agents show performances similar to those of DA markets with ZI agents, in terms of traded quantity, price dispersion and market efficiency when the market size increases. We interpret this as further confirmation that it is agents' strategic sophistication the main cause of the OTC markets' efficiency failure. In fact, as our model predicts, under incomplete information on buyers' valuations and sellers' costs, an increase in market size leads to two counteracting effects: agents' acceptance rates in earlier periods decrease, and earlier offers increase, but the second effect is not always enough to compensate the decrease in acceptance rates.

The main limitation of our study relies on the methodology of classroom experiments. In fact, due to the large subject pool and to the fact that participants in our experiments could not be paid (each of the 300 market sessions took place before tutorials of the first-year introductory course in Microeconomics), we were constrained not to use monetary incentives. We are aware that there are some experimental studies questioning the issue of whether monetary incentives are really necessary to motivate experimental subjects (see, e.g., Holt, 1999, Guala, 2005, and Bardsley et al., 2009).[10] However, we acknowledge that monetary incentives are important in market experiments. Despite that, we stress the point

[10] For example, Camerer and Hogarth (1999) reviewed 74 experiments with no, low, or high performance-based monetary incentives, and found that the modal result has no effect on mean performance.

that our study is comparative: behavior in OTC is analyzed in contrast to behavior in DA and behavior of humans in OTC is analyzed in contrast to behavior of ZI in OTC. Hence, the absence of monetary incentives should not affect our main comparative results.

Our theory-driven experimental study could be easily extended to analyze other specific features of OTC markets that have been previously analyzed in other markets, but not yet in OTC. For example, one may check whether our results hold with asymmetric number of buyers and sellers. Testing OTC trading institutions in markets with few sellers might help understand whether the OTC bargaining rules are able to discipline monopolistic behavior, a feature that is not shared by DA trading mechanisms (see, e.g., Muller et al., 2002).

Another interesting extension could be to test whether the experimentally detected DA-OTC efficiency gap holds under different ratios of intra-marginal over extra-marginal agents. Recall that all our experimental sessions were characterized by a 80-20 fraction of intra-marginal over extra-marginal agents. We find that under this condition prices in the OTC market with human agents are generally lower than the competitive equilibrium price, mainly because of a persistent selling pressure: when public information about existing bids and asks is not available, sellers feel much more pressure than buyers to find a trading counterpart. This in turn leaves room for positive surplus for buyers who would be excluded from transactions if the competitive equilibrium were obtained. Here our intuition is that with a mirrored ratio (20-80) of intra-marginal over extra-marginal agents, the few intra-marginal buyers would also feel a sort of (buying) pressure, thereby disclosing more information about their redemption values to sellers. This could ultimately lead to an increase in the trading prices, which might better converge to the equilibrium, eventually leading to full efficiency. In that case a regulator protecting (or compensating) intra-marginal sellers would not be needed: agents' strategic behavior itself would provide the necessary level of transparency that OTC market institutions need in order to fill the efficiency gap with more centralized trading mechanisms.

References

Adam B. Ashcraft and Darrell Duffie. Systemic Illiquidity in the Federal Funds Market. *American Economic Review*, 97(2):221–225, 2007.

Giuseppe Attanasi, Samuele Centorrino, and Ivan Moscati. Over-the-counter markets vs. double auctions: A comparative experimental study. *Journal of Behavioral and Experimental Economics*, 63:22–35, 2016.

Giuseppe Attanasi, Kene Boun My, Andrea Guido, and Mathieu Lefevbre. Controlling monopoly power in a classroom double-auction market experiment. Working Paper No. 2019-08, Bureau d'Economie Théorique et Appliquée, UDS, Strasbourg, 2019.

Nicholas Bardsley, Robin Cubitt, Graham Loomes, Peter Moffatt, Chris Starmer, and Robert Sugden. *Experimental Economics: Rethinking the Rules*. Princeton, NJ: Princeton University Press, 2009.

Colin Camerer and Robin Hogarth. The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework. *Journal of Risk and Uncertainty*, 19 (1):7–42, 1999.

Timothy N. Cason and Daniel Friedman. Price Formation in Double Auction Markets. *Journal of Economic Dynamics and Control*, 20:1307–1337, 1996.

Timothy N. Cason and Daniel Friedman. A Comparison of Market Institutions. In Charles R. Plott and Vernon L. Smith, editors, *Handbook of Experimental Economics Results*, volume 1, pages 264–272. Amsterdam - New York: North Holland, 2008.

Edward H. Chamberlin. An Experimental Imperfect Market. *The Journal of Political Economy*, 56(2):95–108, 1948.

Kalyan Chatterjee and William Samuelson. Bargaining under complete information. *Operations Research*, 93:835–851, 1983.

Martin W Cripps and Jeroen M Swinkels. Efficiency of large double auctions. *Econometrica*, 74(1):47–92, 2006.

Darrell Duffie. Asset Price Dynamics with Slow-Moving Capital. *The Journal of Finance*, 65(4):1237–1267, 2010.

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

Darrell Duffie and Gustavo Manso. Information Percolation in Large Markets. *American Economic Review*, 97(2):203–209, 2007.

Darrell Duffie, Nicolae Gârleanu, and Lasse Heje Pedersen. Over-the-Counter Markets. *Econometrica*, 73(6):1815–1847, 2005.

Darrell Duffie, Nicolae Gârleanu, and Lasse Heje Pedersen. Valuation in Over-the-Counter Markets. *Review of Financial Studies*, 20(6):1865–1900, 2007.

Darrell Duffie, Semyon Malamud, and Gustavo Manso. Information Percolation With Equilibrium Search Dynamics. *Econometrica*, 77(5):1513–1574, 2009.

Darrell Duffie, Gaston Giroux, and Gustavo Manso. Information Percolation. *American Economic Journal: Microeconomics*, 2(1):100–111, 2010a.

Darrell Duffie, Semyon Malamud, and Gustavo Manso. The Relative Contributions of Private Information Sharing and Public Information Releases to Information Aggregation. *Journal of Economic Theory*, 145(4):1574–1601, 2010b.

Darrell Duffie, Semyon Malamud, and Gustavo Manso. Information Percolation in Segmented Markets. *Journal of Economic Theory*, 153:1–32, 2014.

Urs Fischbacher. z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics*, 10(2):171–178, 2007.

Daniel Friedman. *The Double Auction Market: Institutions, Theories, and Evidence*. Routledge, 2018.

Daniel Friedman and John P. Rust, editors. *The Double Auction Market: Institutions, Theories, and Evidence*. Reading, MA: Addison-Wesley, 1993.

Dhananjay Dan K Gode and Shyam Sunder. Double auction dynamics: structural effects of non-binding price controls. *Journal of Economic Dynamics and Control*, 28(9):1707–1731, 2004.

Dhananjay K Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.

Dhananjay K Gode and Shyam Sunder. What makes markets allocationally efficient? *The Quarterly Journal of Economics*, 112(2):603–630, 1997.

Dhananjay K Gode and Shyam Sunder. Lower bounds for efficiency of surplus extraction in double auctions. In *The Double Auction Market*, pages 199–220. Routledge, 2018.

Thomas A Gresik and Mark A Satterthwaite. The rate at which a simple market converges to efficiency as the number of traders increases: An asymptotic result for optimal trading mechanisms. *Journal of Economic Theory*, 48(1):304–332, 1989.

Francesco Guala. *The Methodology of Experimental Economics*. Cambridge, UK: Cambridge University Press, 2005.

Charles A. Holt. Classroom Games: Trading in a Pit Market. *The Journal of Economic Perspectives*, 10(1):193–203, 1996.

Charles A. Holt. Teaching Economics with Classroom Experiments. *Southern Economic Journal*, 65(3):603–610, 1999.

John A. List. Testing Neoclassical Competitive Market Theory in the Field. *Proceedings of the National Academy of Sciences*, 99(24):15827–15830, 2002.

John A. List. Testing Neoclassical Competitive Theory in Multilateral Decentralized Markets. *Journal of Political Economy*, 112(5):1131–1156, 2004.

R Andrew Muller, Stuart Mestelman, John Spraggon, and Rob Godby. Can double auctions control monopoly and monopsony power in emissions trading markets? *Journal of environmental economics and management*, 44(1):70–92, 2002.

Charles R. Plott. Properties of Disequilibrium Adjustment in Double Auction Markets. In Charles R. Plott and Vernon L. Smith, editors, *Handbook of Experimental Economics Results*, volume 1, pages 16 – 21. Amsterdam - New York: North Holland, 2008.

John Rust, John H Miller, and Richard Palmer. Behavior of trading automata in a computerized double auction market. In *The Double Auction Market*, pages 155–198. Routledge, 2018.

Mark A Satterthwaite and Steven R Williams. The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. *The Review of Economic Studies*, 56(4):477–498, 1989.

Vernon L. Smith. An Experimental Study of Competitive Market Behavior. *The Journal of Political Economy*, 70(2):111–137, 1962.

Robert Wilson. Incentive efficiency of double auctions. *Econometrica*, pages 1101–1115, 1985.

Robert Wilson. Equilibrium in bid-ask markets, g. feiwel, ed. *Arrow and the Ascent of Economic Theory*, 1987.

Appendix

Appendix A: Proof of Proposition 1

Second period: acceptance strategies. In the second period, if an agent receives an offer, he/she accepts it, provided that the offer induces a non-negative payoff. Therefore the acceptance strategies are:

$$s_2^\beta = \begin{cases} \text{Yes} & \text{if } p_2 \leq \beta \\ \text{No} & \text{otherwise;} \end{cases} \qquad s_2^\sigma = \begin{cases} \text{Yes} & \text{if } p_2 \geq \sigma \\ \text{No} & \text{otherwise.} \end{cases}$$

Consider now the offer stage of the second period.

Second period: buyers' offers. Let us first analyze the buyers' optimal strategies, from the lowest to the highest valuation.

- A buyer with valuation $\beta = 0$ knows that he can only trade if he meets the seller with cost 0 and he proposes $p_2 = 0$ (which the seller accepts given the acceptance strategies described above). Hence, $p_2(\beta = 0) = 0$.

- A buyer with valuation $\beta = 1$ knows that he can trade only with sellers of type $\sigma \in \{0, 1\}$. If he proposes $p_2 = 1$ he makes zero profit, hence he proposes $p_2(\beta = 1) = 0$ and makes expected profit $\mathbb{P}_2[\sigma = 0]$.

- A buyer with valuation $\beta = 2$, knows he can trade with sellers of type $\sigma \in \{0, 1, 2\}$. If he proposes $p_2 = 2$ he makes zero profit. If he proposes $p_2 = 1$ he only trades with sellers of type $\sigma \in \{0, 1\}$ and makes expected profits $\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1]$. If he proposes $p_2 = 0$ he only trades with sellers of type $\sigma = 0$ and makes profit 2 when doing so, hence his expected profit is $2 \cdot \mathbb{P}[\sigma = 0]$. Therefore, he will offer

$$p_2(\beta = 2) = \begin{cases} 0 & \text{if } \mathbb{P}_2[\sigma = 0] > \mathbb{P}_2[\sigma = 1] \\ 1 & \text{otherwise.} \end{cases}$$

- A buyer with valuation $\beta = 3$ knows he can optimally trade with sellers of type $\sigma \in \{0, 1, 2, 3\}$. If he offers $p_2 = 3$ he makes zero profits; if he offers $p_2 = 2$ he makes an expected profit of $\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1] + \mathbb{P}_2[\sigma = 2]$; if he offers $p_2 = 1$ he makes an expected profit of $2\,(\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1])$; finally, if he offers a price $p_2 = 0$ he makes an expected profit of $3 \cdot \mathbb{P}_2[\sigma = 0]$. Therefore, he will offer

$$p_2(\beta = 3) = \begin{cases} 0 & \text{if } \mathbb{P}_2[\sigma = 1] < \min\left\{\frac{\mathbb{P}_2[\sigma=0]}{2}, 2\mathbb{P}_2[\sigma = 0] - \mathbb{P}_2[\sigma = 2]\right\} \\ 1 & \text{if } \mathbb{P}_2[\sigma = 1] > \max\left\{\frac{\mathbb{P}_2[\sigma=0]}{2}, \mathbb{P}_2[\sigma = 2] - \mathbb{P}_2[\sigma = 0]\right\} \\ 2 & \text{otherwise.} \end{cases}$$

- A buyer with valuation $\beta = 4$ knows he can optimally trade with all sellers. If he offers $p_2 = 4$ he makes zero profits; if he offers $p_2 = 3$ he makes an expected profit of $\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1] + \mathbb{P}_2[\sigma = 2] + \mathbb{P}_2[\sigma = 3]$; if he offers $p_2 = 2$ he makes an expected profit of $2\left(\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1] + \mathbb{P}_2[\sigma = 2]\right)$; if he offers $p_2 = 1$ he makes an expected profit of $3\left(\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1]\right)$; finally, if he offers a price $p_2 = 0$ he makes an expected profit of $4 \cdot \mathbb{P}_2[\sigma = 0]$. Therefore, he will offer

$$p_2(\beta = 4) = \begin{cases} 0 & \text{if } \mathbb{P}_2[\sigma = 0] > \max\left\{3\mathbb{P}_2[\sigma = 1], \mathbb{P}_2[\sigma = 1] + \mathbb{P}_2[\sigma = 2], \frac{\mathbb{P}_2[\sigma=1]+\mathbb{P}_2[\sigma=2]+\mathbb{P}_2[\sigma=3]}{3}\right\} \\ 1 & \text{if } \mathbb{P}_2[\sigma = 1] > \max\left\{\frac{\mathbb{P}_2[\sigma=0]}{3}, 2\mathbb{P}_2[\sigma = 2] - \mathbb{P}_2[\sigma = 0], \frac{\mathbb{P}_2[\sigma=2]+\mathbb{P}_2[\sigma=3]-2\mathbb{P}_2[\sigma=0]}{2}\right\} \\ 2 & \text{if } \mathbb{P}_2[\sigma = 2] > \max\left\{\begin{array}{l} \mathbb{P}_2[\sigma = 0] - \mathbb{P}_2[\sigma = 1], \frac{\mathbb{P}_2[\sigma=0]+\mathbb{P}_2[\sigma=1]}{2}, \\ \mathbb{P}_2[\sigma = 3] - \mathbb{P}_2[\sigma = 0] - \mathbb{P}_2[\sigma = 1] \end{array}\right\} \\ 3 & \text{otherwise.} \end{cases}$$

Second period: sellers' strategies. We can symmetrically derive the optimal offers for the sellers, from the highest to the lowest cost.

- A seller with cost $\sigma = 4$ proposes $p_2(\sigma = 4) = 4$ and makes zero profits.
- A seller with cost $\sigma = 3$ proposes $p_2(\sigma = 3) = 4$ and makes expected profit $\mathbb{P}_2[\beta = 4]$.
- A seller with cost $\sigma = 2$ proposes

$$p_2(\sigma = 2) = \begin{cases} 4 & \text{if } \mathbb{P}_2[\beta = 4] > \mathbb{P}_2[\beta = 3] \\ 3 & \text{otherwise.} \end{cases}$$

- A seller with cost $\sigma = 1$ proposes

$$p_2(\sigma = 1) = \begin{cases} 4 & \text{if } \mathbb{P}_2[\beta = 3] < \min\left\{\frac{\mathbb{P}_2[\beta=4]}{2}, 2\mathbb{P}_2[\beta = 4] - \mathbb{P}_2[\beta = 2]\right\} \\ 3 & \text{if } \mathbb{P}_2[\beta = 3] > \max\left\{\frac{\mathbb{P}_2[\beta=4]}{2}, \mathbb{P}_2[\beta = 2] - \mathbb{P}_2[\beta = 4]\right\} \\ 2 & \text{otherwise.} \end{cases}$$

- A seller with cost $\sigma = 0$ proposes

$$p_2(\sigma=0) = \begin{cases} 4 & \text{if } \mathbb{P}_2[\beta=4] > \max\left\{3\mathbb{P}_2[\beta=3], \mathbb{P}_2[\beta=3]+\mathbb{P}_2[\beta=2], \frac{\mathbb{P}_2[\beta=3]+\mathbb{P}_2[\beta=2]+\mathbb{P}_2[\beta=1]}{3}\right\} \\ 3 & \text{if } \mathbb{P}_2[\sigma=1] > \max\left\{\frac{\mathbb{P}_2[\sigma=0]}{3}, 2\mathbb{P}_2[\sigma=2]-\mathbb{P}_2[\sigma=0], \frac{\mathbb{P}_2[\sigma=2]+\mathbb{P}_2[\sigma=3]-2\mathbb{P}_2[\sigma=0]}{2}\right\} \\ 2 & \text{if } \mathbb{P}_2[\beta=2] > \max\left\{\begin{array}{l} \mathbb{P}_2[\beta=4]-\mathbb{P}_2[\beta=3], \frac{\mathbb{P}_2[\beta=4]+\mathbb{P}_2[\beta=3]}{2}, \\ \mathbb{P}_2[\beta=1]-\mathbb{P}_2[\beta=4]-\mathbb{P}_2[\beta=3] \end{array}\right\} \\ 1 & \text{otherwise.} \end{cases}$$

First period: acceptance strategies. Note that the second-period expected payoff of types $\beta = 0$ and $\sigma = 4$ is zero, the expected payoff of all other types is positive but (weakly) smaller than 1. We show this for buyers, and the result can be derived symmetrically for sellers. In order to compute this, recall that the expected payoff from period 2 is the average between the expected payoff of making an offer, and the expected payoff of receiving an offer.

- A buyer of type $\beta = 0$ has zero payoff.

- A buyer of type $\beta = 1$ has expected payoff from making an offer equal to $\mathbb{P}_2[\sigma = 0] < 1$, and expected payoff from receiving an offer smaller than 1, so that the average is smaller than 1.

- A buyer of type $\beta = 2$ has expected payoff from making an offer $\mathbb{P}_2[\sigma = 0]+\max\{\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1]\}$, which is always (weakly) smaller than 1, as sellers of type $\sigma \in \{3, 4\}$ never trade in the proposed equilibrium, and therefore both $\mathbb{P}_2[\sigma = 0]$ and $\mathbb{P}_2[\sigma = 1]$ are at most $\frac{1}{3}$; the expected payoff from receiving an offer is also smaller than 1, so that the average is smaller than 1.

- A buyer of type $\beta = 3$ has expected payoff from making an offer either equal to $\sum_{k=0}^{2} \mathbb{P}[\sigma = k] < 1$, or $2\left(\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1]\right)$, which is smaller than 1 because sellers of type $\sigma \in \{3, 4\}$ never trade in the proposed equilibrium, and therefore both $\mathbb{P}_2[\sigma = 0]$ and $\mathbb{P}_2[\sigma = 1]$ are at most $\frac{1}{3}$; or $3 \cdot \mathbb{P}_2[\sigma = 0] < 1$ for the same reason. The expected payoff from receiving an offer is also smaller than 1, so that the average is smaller than 1.

- Consider now a buyer of type $\beta = 4$. Recall that in equilibrium sellers of type $\sigma \in \{2, 3, 4\}$ never trade in the first period. As a consequence, choosing $p_2 = 2$ leads to the (weakly) highest payoff of $2\left(\sum_{k=0}^{2} \mathbb{P}[\sigma = k]\right)$, which is (weakly) higher than 1 (but smaller than 2). The maximum value that this expected payoff can take is $\frac{6}{5}$, which happens when all the sellers of type $\sigma \in \{0, 1, 2\}$ are still in the market in period 2. As for the expected payoff from receiving an offer, the maximum expected payoff from period 2 is $\frac{4}{5}$, which happens when all the sellers of type $\sigma \in \{0, 1, 2\}$ are still in the market in period 2. Therefore, the expected continuation payoff of a buyer of type $\beta = 4$ is bounded above by $\frac{1}{2}\left(\frac{6}{5} + \frac{4}{5}\right) = 1$.

Therefore, in the first period, the acceptance strategies are:

$$s_2^\beta = \begin{cases} \text{Yes} & \text{if } p_2 \leq \max\{0, \beta - 1\} \\ \text{No} & \text{otherwise;} \end{cases} \qquad s_2^\sigma = \begin{cases} \text{Yes} & \text{if } p_2 \geq \min\{\sigma + 1, 4\} \\ \text{No} & \text{otherwise.} \end{cases}$$

Consider now the offer stage of the first period.

First period: buyers' offers. Let us first analyze the buyers' optimal strategies, from the lowest to the highest valuation.

- A buyer with valuation $\beta = 0$ knows that he can never trade in the first period. Therefore he proposes $p_1(\beta = 0) = 0$, as any other price may induce a negative payoff.
- A buyer with valuation $\beta = 1$ knows he can trade if he meets a seller with cost $\sigma = 0$. However, his only opportunity to trade with a seller of type $\sigma = 0$ is to offer $p_1 = 1$. In this case, he would forego the possible positive payoffs of period 2, therefore $p_1(\beta = 1) = 0$.
- A buyer with valuation $\beta = 2$ maximizes his/her profits by offering $p_1(\beta = 2) = 1$.
- A buyer with valuation $\beta = 3$ maximises his/her profits by making an offer $p_1(\beta = 3) = 2$.
- A buyer with valuation $\beta = 4$ maximises his/her profits by making an offer $p_1(\beta = 3) = 2$.

First period: sellers' offers. Sellers' offers can be derived symmetrically.

Therefore, buyers and sellers' first-period equilibrium offers are, respectively:

$$p_1^\beta = \begin{cases} 0 & \text{if } \beta \in \{0, 1\} \\ 1 & \text{if } \beta = 2 \\ 2 & \text{if } \beta \in \{3, 4\}, \end{cases} \qquad p_1^\sigma = \begin{cases} 2 & \text{if } \sigma \in \{0, 1\} \\ 3 & \text{if } \sigma = 2 \\ 4 & \text{if } \sigma \in \{3, 4\}. \end{cases}$$

■

Appendix B: The case of $n = 6$ and the effect of market size

Section 3.1 contains the analysis of the market with $n = 10$. Here we analyze the $n = 6$ market and compare the strategic behavior of the two markets. Note that we model the change in market size in a manner which is as close as possible to our experimental conditions (see subsection "Market sizes" in the experimental design of human agents of Section 2.1): the lowest and highest redemption values are unchanged across markets, but the distance between values increases in the smaller market (2 instead of 1).

The case of six traders

We consider a $n = 6$ market with the same structural characteristics as the $n = 10$ market described in Section 3. The market is now composed by 3 buyers and 3 sellers. Buyers privately evaluate the object $v \in \{0, 2, 4\}$, while sellers have a production cost $c \in \{0, 2, 4\}$. As before (and as in our experiment), there is exactly one buyer (seller) with each valuation (cost), so that each valuation (cost) is equally likely.

Proposition 2 *The following is the subgame perfect equilibrium of the two-period bargaining game.*

- *Period t offers, for $t = 1, 2$, are:*

$$p_t^\beta = \begin{cases} 0 & \textit{if } \beta \in \{0,2\} \\ 2 & \textit{if } \beta = 4, \end{cases} \qquad p_t^\sigma = \begin{cases} 2 & \textit{if } \sigma = 0 \\ 4 & \textit{if } \sigma \in \{2,4\}. \end{cases}$$

- *Period 1 acceptance strategies are:*

$$s_2^\beta = \begin{cases} \textit{Yes} & \textit{if } p_2 \leq \max\{0, \beta - 2\} \\ \textit{No} & \textit{otherwise}, \end{cases} \qquad s_2^\sigma = \begin{cases} \textit{Yes} & \textit{if } p_2 \geq \min\{\sigma + 2, 4\} \\ \textit{No} & \textit{otherwise}. \end{cases}$$

- *Period 2 acceptance strategies are:*

$$s_2^\beta = \begin{cases} \textit{Yes} & \textit{if } p_2 \leq \beta \\ \textit{No} & \textit{otherwise;} \end{cases} \qquad s_2^\sigma = \begin{cases} \textit{Yes} & \textit{if } p_2 \geq \sigma \\ \textit{No} & \textit{otherwise}. \end{cases}$$

Proof:

Second period: acceptance strategies. In the second period, if an agent receives an offer, he/she accepts it, provided that the offer induces a non-negative payoff. Therefore the acceptance strategies are:

$$s_2^\beta = \begin{cases} \text{Yes} & \text{if } p_2 \leq \beta \\ \text{No} & \text{otherwise;} \end{cases} \qquad s_2^\sigma = \begin{cases} \text{Yes} & \text{if } p_2 \geq \sigma \\ \text{No} & \text{otherwise}. \end{cases}$$

Consider now the offer stage of the second period.

Second period: buyers' offers. Let us first analyze the buyers' optimal strategies.

- A buyer with valuation $\beta = 0$ knows that he can only trade if he meets the seller with cost 0 and he proposes $p_2 = 0$ (which the seller accepts given the acceptance strategies described above). Hence, $p_2(\beta = 0) = 0$.

- A buyer with valuation $\beta = 2$ knows that he can trade only with sellers of type $\sigma \in \{0, 2\}$. If he proposes $p_2 = 2$ he makes zero profit, hence he proposes $p_2(\beta = 2) = 0$ and makes expected profit $2 \cdot \mathbb{P}_2[\sigma = 0]$.

- A buyer with valuation $\beta = 4$ knows he can optimally trade with all sellers. If he offers $p_2 = 4$ he makes zero profits; if he offers $p_2 = 2$ he makes an expected profit of $2\,(\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 2])$; finally, if he offers a price $p_2 = 0$ he makes an expected profit of $4 \cdot \mathbb{P}_2[\sigma = 0]$. Therefore, he will offer

$$p_2(\beta = 4) = \begin{cases} 0 & \text{if } \mathbb{P}_2[\sigma = 0] > \mathbb{P}_2[\sigma = 1] \\ 2 & \text{otherwise.} \end{cases}$$

Second period: sellers' strategies. We can symmetrically derive the following optimal offers for the sellers.

- A seller with cost $\sigma = 4$ proposes $p_2(\sigma = 4) = 4$ and makes zero profits.

- A seller with cost $\sigma = 2$ proposes $p_2(\sigma = 2) = 4$ and makes expected profit $2 \cdot \mathbb{P}_2[\beta = 4]$.

- A seller with cost $\sigma = 0$ proposes

$$p_2(\sigma = 2) = \begin{cases} 4 & \text{if } \mathbb{P}_2[\beta = 4] > \mathbb{P}_2[\beta = 2] \\ 2 & \text{otherwise.} \end{cases}$$

First period: acceptance strategies. Note that the second-period expected payoff of types $\beta = 0$ and $\sigma = 4$ is zero, the expected payoff of all other types is positive but (weakly) smaller than 2. We show this for buyers, and the result can be derived symmetrically for sellers. In order to compute this, recall that the expected payoff from period 2 is the average between the expected payoff of making an offer, and the expected payoff of receiving an offer.

- A buyer of type $\beta = 0$ has zero payoff.

- A buyer of type $\beta = 2$ has expected payoff from making an offer equal to $2 \cdot \mathbb{P}_2[\sigma = 0] < 2$, as $\mathbb{P}_2[\sigma = 0] < \frac{1}{2}$, because the seller of type $\sigma = 4$ never trades in the first period. Moreover, the expected payoff from receiving an offer is smaller than 2, so that the average is smaller than 2.

- A buyer of type $\beta = 4$ has expected payoff from making an offer

$$2\left(\mathbb{P}_2[\sigma = 0] + \max\{\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1]\}\right)$$

  which is always smaller than 2, as the seller of type $\sigma = 4$ never trades in the proposed equilibrium, and therefore both $\mathbb{P}_2[\sigma = 0]$ and $\mathbb{P}_2[\sigma = 1]$ are at most $\frac{1}{2}$, and their sum is at most $\frac{1}{3}$; the expected payoff from receiving an offer is also smaller than 2, so that the average is smaller than 2.

Therefore, in the first period, the acceptance strategies are:

$$s_2^{\beta} = \begin{cases} \text{Yes} & \text{if } p_2 \leq \max\{0, v(\beta) - 2\} \\ \text{No} & \text{otherwise;} \end{cases} \qquad s_2^{\sigma} = \begin{cases} \text{Yes} & \text{if } p_2 \geq \min\{c(\sigma) + 2, 4\} \\ \text{No} & \text{otherwise.} \end{cases}$$

Consider now the offer stage of the first period.

Buyers' first-period offers. Let us first analyze the buyers' optimal offers in the first period, from the lowest to the highest valuation.

- A buyer with low valuation, $\beta = 0$, knows that he can never trade in the first period. Therefore he proposes $p_1(\beta = 0) = 0$, as any other price may induce a negative payoff.

- A buyer with intermediate valuation, $\beta = 2$, knows he can trade if he meets a seller with low cost. However, his only opportunity to trade with a seller of type $\sigma = 0$ is to offer $p_1 = 2$. In this case, he would forego the possible positive payoffs of period 2, therefore $p_1(\beta = 2) = 0$.

- A buyer with high valuation, $\beta = 4$, knows he can optimally trade with any potential seller. However, in order to trade with sellers of type $\sigma \in \{2, 4\}$ he has to offer $p_1 = 4$ which gives zero payoff. Given that his expected payoff from period 2 is positive and smaller than 2, his optimal choice is to offer $p_1 = 2$ which gives him a payoff of 2 in case the offer his accepted (which happens with probability $\mathbb{P}_1[\sigma = 0] = \frac{1}{3}$). Therefore $p_1(\beta = 4) = 2$.

Sellers' first-period offer. We can symmetrically derive the sellers' optimal offers in the first period, from the highest to the lowest cost.

- A seller with high cost, $\sigma = 4$, offers $p_1(\sigma = 4) = 4$, as any other price may induce a negative payoff.

- A seller with intermediate cost, $\sigma = 2$, offers $p_1(\sigma = 2) = 4$.

- A seller with low cost, $\sigma = 0$, offers $p_1(\sigma = 0) = 2$.

Therefore, buyers and sellers' first-period equilibrium offers are, respectively:

$$p_1^\beta = \begin{cases} 2 & \text{if } \beta = 4 \\ 0 & \text{otherwise,} \end{cases} \quad p_1^\sigma = \begin{cases} 2 & \text{if } \sigma = 0 \\ 4 & \text{otherwise.} \end{cases}$$

Note that in the first period trade happens only if a buyer of type $\beta = 4$ meets with a seller of type $\sigma = 0$, regardless of the proposer. Therefore, given that all types are equally likely before trade happens, it is always the case that $\mathbb{P}_2[\beta = 4] \leq \mathbb{P}_2[\beta = 2]$ and that $\mathbb{P}_2[\sigma = 0] \leq \mathbb{P}_2[\sigma = 2]$. As a consequence, $p_2(\beta = 4) = p_2(\sigma = 0) = 2$. ■

Effects of market size: the comparison between $n = 6$ and $n = 10$

If we compare Proposition 1 with Proposition 2 we first notice that, for a given redemption value, traders are less willing to accept less advantageous offers in the first period. Second, traders anticipate this effect and modify first-period offers when the market size increases, so that buyers' offers are weakly higher and sellers' ones are weakly lower when $n = 10$ with respect to the market with $n = 6$.

We now ask ourselves how this translates into changes of the market features that we test experimentally. Table B.1 shows the expected probability of observing, in equilibrium, a number of trades lower, equal or greater than the efficient one (which is 2 for the $n = 6$ market, and 3 for the $n = 10$ market). It also shows the expected value of the efficiency index (expected surplus over efficient surplus, where the efficient surplus is 4 for the $n = 6$ market and 7 for the $n = 10$ market).[11] We notice that the expected probability of having a lower than efficient trading volume increases with market size, and that the expected efficiency decreases. Therefore, the simulations in Table B.1 show that the first strategic effect of a market size increase – i.e., traders being less willing to accept less advantageous offers in period 1 – is stronger then the second one – i.e., buyers' (resp., sellers') offers being weakly higher (resp., lower) in period 1.

	$n = 6$	$n = 10$
$q < q^*$	75.01%	84.83%
$q = q^*$	24.99%	15.13%
$q > q^*$	0.00%	0.04%
Efficiency	77.77%	64.10%

Table B.1: Expected performance of the two-period theoretical OTC market, by size

[11] The expected performance has been computed by means of a Python code which is available upon request from the authors.