

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)

Abstract

We study two well-known electronic markets: an over-the-counter (OTC) market, in which each trader 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-intelligence agent-based simulations. Simulation results show that the traded quantity increases with market size under both DA and OTC. Experimental results confirm the same tendency under DA, while the share of periods in which the traded quantity is lower than the efficient one increases with market size under OTC, ultimately leading to a DA-OTC efficiency gap increasing with the market size. We rationalize these results by putting forward a simple model of OTC market as a repeated bargaining procedure under incomplete information on buyers’ valuations and sellers’ costs. We show that efficiency decreases slightly with size due to two counteracting effects: acceptance rates in earlier interactions decrease with size, and earlier offers increase, but not always enough to compensate for the decrease in acceptance rates.

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

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

Acknowledgements: We thank Dhananjay K. Gode and Shyam Sunder for sharing the code for agent-based simulations of zero-intelligence 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, Ornella Lepre, 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. One of the most common market structures considered by the experimental literature is the double auction (henceforth, DA) mechanism introduced by Vernon Smith (1962). This mechanism has so far been used to test the behavior of competitive markets (Friedman and Rust, 1993; Plott, 2008; Cason and Friedman, 2008) and the impact of public intervention (Cox et al., 2018; Morone et al., 2018), and to prove that zero-intelligence (ZI) agents 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, trading 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; Holt, 1996; 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. However, there are 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 OTC relative performance 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 ascribe 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., Friedman and Rust, 1993; 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. The same occurs under agent-based simulations with ZI agents (for a review, see Gode and Sunder, 2018; Rust et al., 2018). To the best of our knowledge, similar robustness tests to market size modifications are missing for OTC markets in both streams of research, as mentioned above. Therefore, in this paper, we investigate OTC markets through the simultaneous use of experiments with human traders and agent-based simulations with ZI agents, with a threefold aim.

First, we want to understand whether humans may perform as well as ZI agents 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. With this, we build on the work by Attanasi et al. (2016), and depart from it, as they only compare OTC and DA markets with humans for a fixed number of agents (40), without exploring whether the two types of markets, and henceforth their comparison, are affected by 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 ZI agents as a function of the market size. In this analysis, we consider the efficiency loss between humans and ZI agents 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 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. In contrast, 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 DA market where all agents simultaneously submit bids and asks 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 the efficiency of OTC markets is affected by market size is scarce. 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. In addition, most of the strategic models of decentralized markets (see, e.g., Moreno and Wooders, 2002; Blouin and Serrano, 2001; Lauermaun, 2012) do not consider how the behavior of decentralized markets varies with size. An exception is Wooders (1998), who compares equilibrium behavior of small and large decentralized markets, without considering however the efficiency issue. For this reason, we put forward a simple model of OTC markets, which focuses on the strategic effects at play and their effects on the relationship between efficiency and market size. Our game-theoretical model shows that the intuition of Wilson (1987) on why DA markets with a large number of traders become 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.¹ As a consequence, an increase in market size reduces efficiency in OTC markets (see Coleff and Garcia, 2017, for a similar result in procurement auctions with private information).

We run 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 an equal number of buyers and sellers (10, 20, 40, or 80 traders). Due to the massive number of students involved, and as is indeed prevalent in many classroom experiments (see, e.g., Holt, 1996, 1999; Attanasi et al., 2016), we did not use monetary incentives. Instead, 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

¹This finding is in line with a strand of literature on bargaining, which highlights different reasons for strategic delay (see, for example, Deneckere and Liang, 2006; Ausubel et al., 2002). In this literature, the study that is closer in spirit to our work is Bochet and Siegenthaler (2018). They propose a theory-driven experiment on bargaining with finite horizon in the presence of asymmetric information and frictions and investigate how efficiency responds to changes in the (finite) time horizon.

implementing algorithms of zero-intelligence computerized trading *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 efficiency under different market sizes.

Our results on the comparison between OTC markets and DA markets with different types of traders (humans vs. ZI traders) and different market sizes (10, 20, 40 and 80 traders) show that, while for ZI traders the two trading mechanisms respond similarly to changes in market size, for human agents this is not true. In particular, OTC markets with humans respond to market size in a way that substantially differs from DA markets with humans and from DA and OTC markets with ZI agents. 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 traders. 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 ZI agent-based simulations. In Section 3 we present the experimental hypotheses on traded quantity, trading price, efficiency and sources of inefficiency, together with the results of simulations of markets with ZI agents that inform them, and their experimental test. In Section 4 we introduce the theoretical model on OTC markets that explains their behavior (Section 4.1), and finally verify the link between the theoretical model and the experimental results (Section 4.2). 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). Then, we describe the features of the agent-based simulations with ZI agents (Section 2.2). As for human agents, the design of computerized DA and OTC markets represents the electronic version of the corresponding oral DA and OTC markets of the seminal classroom experiments of respectively Smith (1962) and Chamberlin (1948).² As for

²Differently from the oral DA of Smith (1962) and the OTC market of Chamberlin (1948), our markets are fully computerized. Furthermore, as in Smith (1962) and subsequent standard practice in market classroom

ZI agents, we implement the same agent-based simulations of Gode and Sunder (1993) in DA markets and we extend their code to OTC markets under comparable conditions.

2.1 Experimental design with human agents

Procedures. The experimental procedures are essentially the same as in Attanasi et al. (2016), who run their experiments over the six consecutive years 2009 – 2014, while our data come from the subsequent five years 2015 – 2019. All classroom experiments were held at Bocconi University, Milan, during a first-year introductory course in Microeconomics, always in the first two weeks of October (first semester). The experiments were computerized through the z-Tree software (Fischbacher, 2007), run in the same computerized room, and administered by the same experimenter (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). Differently from Attanasi et al. (2016), who performed the DA-OTC comparison under a fixed market size of 40 traders, in our design market size varies. As a matter of fact, each experimental session was characterized by a market mechanism (DA or OTC) and a market size of $n \in \{10, 20, 40, 80\}$ traders ($n/2$ buyers and $n/2$ sellers). The computerized room where each classroom experiment was run had 90 trading positions (PCs), thereby allowing to run at least two experimental sessions simultaneously. 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 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 trading mechanism – market size combination, according to a between-subject design. The 8 market mechanism-size combinations in Table 1 also represent 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 experimental sessions for each market mechanism-size combination (treatment)

experiments (see, e.g., Holt, 1996; Cason and Friedman, 2008), we allow subjects to trade for several periods so as to gain experience about the trading mechanism. Differently, Chamberlin (1948) let subjects trade for only one period. Finally, under both DA (as in Smith, 1962) and OTC (differently from Chamberlin 1948), we always implement post-trade price transparency.

Common features. We first 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 of 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 assignments are kept constant during the whole experiment.³
- *Number and length of trading periods.* Each *session* consists of three phases. Each *phase* is composed of three trading periods of fixed length.⁴ The length of a *period* is 120 seconds in the first two phases, and 60 seconds in the third phase.⁵
- *Tradable units.* In each trading period, each seller (resp., buyer) owns (resp., can purchase) one unit of a homogeneous good, so that each subject can only trade one unit per period, thereby exiting the market after having traded. 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 exogenously assigned a different redemption value for the single unit he/she has to buy or sell. A buyer's redemption value is the maximum amount he/she can spend for one unit of the good, i.e., his/her *valuation*. A seller's redemption value is the minimum amount he/she has to receive for his/her unit, i.e., his/her *cost*.
- *Budget constraints.* 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 complete the trade within the period, he/she has zero profits.
- *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

³We follow Smith (1962) for DA and Chamberlin (1948) for OTC markets in assigning subjects to one of the two roles exogenously and randomly. As for experiments on OTC markets where subjects choose their role as buyer or seller endogenously, see List (2002, 2004).

⁴Note that, although there is not a per period time constraint in the pioneering studies of Smith (1962) and Chamberlin (1948), a per period time limit has later become a quite standard feature of both DA and OTC classroom experiments: see Wells (1991) for DA and Holt (1996) and Ruffle (2003) for OTC. This is a necessary feature in order to experimentally allow intramarginal inefficiency, which has been shown to be a relevant source of inefficiency in electronic markets (see, e.g., Cason and Friedman, 1996).

⁵We took the per period time limit of 110 seconds of electronic markets with $n = 8$ traders in Cason and Friedman (1996) as reference for our 120 seconds time limit for markets with $n = 10$ traders, and maintained it for the other three greater market sizes in Table 1.

valuations and costs in the market.⁶ 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.* Individual profits are shown at the end of each trading period as the difference between valuation and trading price for buyers and between trading price and cost for sellers. 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 phase, the sum of the 3-period profits is corrected so as to compensate for unlucky draws in redemption values. Subjects are then ranked according to their corrected phase profit.⁷ 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.⁸

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

- *DA markets.* Their three main features are: (i) *public trading*: bids and asks are publicly posted, so that at any point in time every subject knows the current highest bid and lowest ask; (ii) *bid/ask improvement rule*: subjects are only allowed to make offers that improve on the current situation, i.e., higher bids or lower asks; (iii) *public information about trading prices*: when a trade is closed, the trading price appears on every subject's screen.

⁶Note 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 numbers as bids and asks.

⁷The correction factor is 1 for the two traders with the best possible redemption values (i.e., maximum valuation and minimum cost) and, for all other traders, is proportionally increasing with the distance between the individual redemption value and the best possible one.

⁸Note that even though classroom experiments typically raise issues of post-experimental communication, we think that in our experiment spillovers between classrooms hardly arose for three reasons, that we share with Attanasi et al. (2016). First, because of the specific features of the subject pool: Bocconi University enrolls a quite heterogeneous population of students in terms of geographical background. Second, because of the timetable of the experiment: all subjects were first-year undergraduate students participating in the experiment during their first month of classes, students of the same class group participated in the same time slot, and within an academic year all time slots were allocated within one week only, with 4-5 classroom experiments per day. Finally, because of experimenter control of the sequence of implemented treatments: the market mechanism-size combination of two subsequent classroom experiments was never the same.

- *OTC markets.* Their three corresponding features are: (i) *private trading*: each subject can make only one offer at a time, by indicating its amount and the counterpart’s ID. This offer can be observed only by the selected counterpart, so that subjects do not observe the highest bid and lowest ask standing in the market; (ii) *no bid/ask improvement rule*: there are no restrictions on the amount that a subject can offer, either in relation to his/her previous offers or to other subjects’ ones; (iii) *public information about trading prices*: if an offer is accepted, the trading price appears on every subject’s screen. Note that the last feature is the same as in DA markets,⁹ while features (i)-(ii) constitute the treatment difference. Design feature (i) deserves a more thorough discussion since there are many possible alternatives to implement private trading in OTC markets (e.g., random vs. endogenous interactions, fixed vs. endogenous duration of each interaction). We opted for endogenously generated interactions in order to match what occurs in oral classroom (Chamberlin, 1948; Holt, 1996; Ruffle, 2003) and field experiments (List, 2002, 2004) and in general in real OTC trading interactions (Hendershott and Madhavan, 2015), where buyer-seller interactions are the outcome of traders’ searches of the best counterpart by milling around the marketplace.¹⁰ For the same reasons, we also opted for an endogenous duration of a buyer-seller interaction, so that a buyer (resp., a seller) is allowed to withdraw his/her bid (resp., ask) by making the offer disappear from the screen of the trading counterpart at any time during the trading period. Therefore, an agent can quickly approach all the agents on the other market side. However, each time he/she approaches a new agent he/she has to withdraw the offer made to the previous one, i.e., he/she can only approach one agent per time.¹¹ As in the corresponding oral markets, it is possible that a subject is approached by more than one agent at the same time, hence getting multiple offers. In this case, he/she sees them already in order, with the best one appearing at the top of his/her screen.

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

⁹Given public information about trading prices, it is possible that subjects’ decisions are anchored on earlier prices in the period, and this may affect inefficiencies.

¹⁰Hendershott and Madhavan (2015) study traditional OTC trading based on telephone and voice communications. In particular, they use data on corporate bond trades between 2010 and 2011 to investigate which factors influence the transition from voice-based OTC trading to DA trading based on electronic platforms such as MarketAxess.

¹¹Note that, given the time limit of 120 seconds per period, a fixed time duration for each offer would have restricted the set of agents’ bargaining strategies in a way that does not apply to DA, where each agent can make as many as (public) offers he/she likes. More importantly, for the specific goal of our study, this exogenous restriction would have been differently binding for different market sizes, thus providing an additional mechanism through which size affects market behavior.

(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 , a size-dependent efficient quantity $q^* = 0.8 \cdot (n/2)$, i.e., a size-independent ratio of efficient over total quantity (80% of the available units are traded in equilibrium), and a size-independent average efficient surplus.

2.2 Design of simulations with ZI agents

We implement agent-based simulations with two groups of zero-intelligence agents, buyers and sellers, by imposing the same market rules as in the experiments with human agents. To increase humans-ZI agents comparability, we made these agent-based simulations by modeling ZI agents within the same z-Tree environment (Fischbacher, 2007) used to run the classroom experiments described in Section 2.1.¹²

Common features. 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 own 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 human agents is that – given that ZI agents have neither memory nor learning – they do not care about information. Hence, 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 for each treatment. 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 and OTC with $n = 10$), 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 ZI trading periods for each market mechanism-size combination (treatment)

¹²Note that this is not part of the codes we adapted from Attanasi et al. (2016), as they did not implement agent-based simulations.

Market mechanisms. In DA markets, agents (randomly) choose only either their ask or bid, entering the market sequentially. In OTC markets, instead, agents (randomly) choose both their ask or bid and the counterpart to whom that ask or bid is going to be sent. In OTC markets, as asks and bids are only received by a single counterpart, agents are allowed to enter the market simultaneously in order to avoid slowing down transactions artificially. In Appendix A 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. Second, we quadruplicate negotiation time whenever the market size is doubled because the number of possible buyer-seller pairs grows by a factor of 4 when the market size doubles. Third, keeping the time length of the trading period fixed across DA and OTC for the same market size provides a measure of OTC intra-marginal inefficiency as compared to DA, i.e., only due to market rules and not to (human) agents' strategic behavior. Indeed, OTC markets have greater complexity than DA markets for any market size, since in the former each buyer (resp., seller) must associate the number of the selected seller (resp., buyer) to his/her private bid (resp., ask). With this, the increase in the number of possible buyer-seller pairs due to greater market size is by construction more relevant for OTC than for DA markets.

3 Simulation and experimental results

This section compares the performance of DA and OTC mechanisms by market size in terms of traded quantity (Section 3.1), trading price (Section 3.2), efficiency (Section 3.3), and sources of inefficiency (Section 3.4).

Each section both puts forward the hypothesis and tests it. In order to motivate our hypotheses on the effects of market mechanism and size on human agents' behavior, we integrate the known findings in the theoretical and experimental literature with our agent-based simulations. More precisely, we first indicate the theme of interest and describe the theoretical and experimental findings in the literature on markets with humans on that theme. Second, for the same theme, we report the behavior of markets with ZI agents in our simulations. The latter are useful to elaborate hypotheses when neither the theoretical

nor experimental literature on markets with humans has yet provided findings. Then, we formulate the hypothesis.

Each hypothesis is divided into three parts. First, it focuses on the trend of human agents' behavior as market size increases from 10 to 80 agents under DA. Second, it considers the corresponding trend under OTC. Finally, it assesses the trend of the DA-OTC gap as market size increases.

Note that the theoretical and experimental literature on markets with humans has produced thorough findings as for size trends in DA, and, for fixed market size, also for the DA-OTC gap. Therefore, our agent-based simulations are mainly meant to develop hypotheses about size trends in OTC markets, thereby leading us to integrate previous findings on DA markets in order to elaborate informed predictions on the size trend of the DA-OTC gap.

In the analysis of the results, we first test each part of the hypothesis and then perform the comparison between ZI and human agents within each market structure. This helps shading light on the interplay of humans' strategic behavior and market size, which is our main focus of interest.

3.1 Traded quantity

The experimental literature on human agents suggests that volumes of trade in DA markets tend to the efficient quantity when market size increases (Friedman and Rust, 1993; Friedman, 2018), and that DA markets induce higher volumes of trade than OTC for markets of size 40 (Attanasi et al., 2016).

The left panels of Tables 3-4 report the results of simulations with ZI agents. Table 3 shows the percentage of periods in which the traded quantity is lower, equal, or higher than the efficient one, by trading mechanism and market size. Table 4 reports the results of probit regressions where the dependent variables indicate respectively whether the traded quantity is below and above the efficient one. The explanatory variables are a mechanism dummy for OTC markets (with DA as baseline), three dummies for market size (with size 10 as baseline), and the interaction between the mechanism dummy and the market size dummies. Recall that, in our agent-based simulations, we set the length of trading periods increasing with the market size so as to compensate for increased complexity of interactions among ZI agents. With this, Table 3 reports that, in line with previous results of simulations for ZI agents under DA (for a review, see Gode and Sunder, 2018; Rust et al., 2018), the traded quantity under DA is higher than the efficient one more often as market size increases.¹³

¹³In our DA markets with ZI agents, the traded quantity is never lower than the efficient one. This is a combination of efficiency of DA markets with ZI agents and the fact that we chose the length of the trading period for DA markets of size $n = 10$ with the specific aim of allowing enough time for ZI traders to reach the efficient quantity (see Gode and Sunder, 1993), at the same time increasing this length for higher size n proportionally to the increased complexity of ZI agents' interactions. In this way, the DA market of size 10

	ZI agents			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 3: Percentage of periods in which the traded quantity is lower, equal or higher than the efficient one, $q^* = 0.8 \cdot (n/2)$, by type of traders, treatment and market size $n \in \{10, 20, 40, 80\}$

	ZI Agents		Humans	
	$q < q^*$	$q > q^*$	$q < q^*$	$q > q^*$
OTC	3.276 (0.981)	-0.307 (0.000)	0.580 (0.000)	-0.246 (0.003)
20	0.000 (1.000)	0.729 (0.000)	-0.121 (0.152)	0.221 (0.016)
40	0.000 (1.000)	1.340 (0.000)	-0.422 (0.000)	0.995 (0.000)
80	0.000 (1.000)	1.699 (0.000)	-0.691 (0.000)	1.659 (0.000)
OTC \times 20	0.802 (0.997)	-0.544 (0.000)	-0.081 (0.487)	0.208 (0.117)
OTC \times 40	1.340 (0.995)	-1.008 (0.000)	0.443 (0.004)	-0.663 (0.000)
OTC \times 80	1.242 (0.995)	-1.043 (0.000)	1.641 (0.000)	-2.039 (0.000)
Const.	6.083 (0.966)	-0.471 (0.000)	-0.419 (0.000)	-0.967 (0.000)

Table 4: Probit regression of a dummy for quantity below and above the efficient one, $q^* = 0.8 \cdot (n/2)$, over trading mechanism and market size $n \in \{10, 20, 40, 80\}$ (p -values in brackets)

We verify that this also occurs under OTC. Indeed, also for OTC markets Table 3 shows that the traded quantity is usually higher than the efficient one (χ^2 test for $q > q^*$ vs. $q \leq q^*$, p -value < 0.001 for each market size), and that the percentage of periods with a traded quantity higher than the efficient one significantly increases with market size (+23 percentage points moving from size 10 to size 80, this increase being significant at the 0.1%

works as a baseline for the performance of DA markets of larger size and of OTC markets of any size (see subsection “Market sizes” in the design of simulations with ZI agents of Section 2.2).

level, χ^2 test).

However, we find significantly lower traded quantity under OTC than under DA regardless of the market size, since the negative coefficient of the OTC dummy for $q > q^*$ in Table 4 is significant at the 1% level. This is in line with the results of Attanasi et al. (2016) for human agents with $n = 40$. We interpret our simulation results as confirmation that the comparatively greater complexity of the OTC mechanism with respect to the DA mechanism lowers traded quantity regardless of market size. Furthermore, the DA-OTC gap in the traded quantity increases with market size, since the negative coefficients of dummies OTC \times 20, OTC \times 40 and OTC \times 80 for $q > q^*$ are significant and increasingly larger.¹⁴

All this allows us to state hypothesis **H1**, which summarizes our predictions on traded quantity for human agents.

H1: (DA) In DA markets the traded quantity is less often lower than the efficient one as market size increases;

(OTC) In OTC markets the traded quantity is less often lower than the efficient one as market size increases;

(DA-OTC) OTC markets have lower traded quantity than DA markets for each market size, and the DA-OTC gap in traded quantity increases with market size.

Let us now test **H1**. The right panels of Table 3 and 4 report the results for experimental markets with humans.

DA markets respond to market size as in the case with ZI agents – see the right panel of Table 3, especially column $q > q^*$ for DA (+49 percentage points moving from size 10 to size 80, increase significant at the 0.1% level, χ^2 test). Therefore, **H1(DA) is confirmed**.

OTC markets instead display the opposite behavior than DA markets – see the right panel of Table 3, especially column $q < q^*$ for OTC (+30 points moving from size 10 to size 80, increase significant at the 0.1% level, χ^2 test). Therefore, **H1(OTC) is rejected**.

As for the DA-OTC gap, the OTC coefficient in the right panel of Table 4 is positive (resp., negative) and significant for $q < q^*$ (resp., for $q > q^*$), thus confirming that the traded quantity is more likely to be below the efficient one in OTC markets than in DA ones. Moreover, the coefficients of OTC \times 40 e OTC \times 80 are positive (resp., negative) and significant for $q < q^*$ (resp., $q > q^*$). Thus, the likelihood of observing a traded quantity lower than the efficient one increases in OTC markets with respect to DA markets for sizes $n \geq 40$. Hence, **H1(DA-OTC) is confirmed**.¹⁵

¹⁴Note that this negative mechanism effect is only of second order with respect to the positive size effect: 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).

¹⁵Note that we formulated H1(DA-OTC) by relying on the assumption that traded quantity in OTC

ZI-human agents comparison. Comparing the right and left panels of Table 3, we see that under OTC the human-ZI agents (resp., ZI-human agents) difference in the probability of observing $q < q^*$ (resp., $q > q^*$) increases monotonically and significantly from 56 (resp., 45) percentage points for $n = 10$ markets to 81 (resp., 74) percentage points for $n = 80$ markets (t-test of the difference in percentages, $p\text{-value} < 0.01$). Therefore, for OTC markets, the ZI-human agents gap in terms of traded quantity is increasing in market size. The control for DA markets shows the opposite picture, with the ZI-human agents gap in terms of traded quantity decreasing in market size.¹⁶

To summarize, the test of H1 and the ZI-human agents comparison show that in DA markets humans' size-dependent behavior is similar to the one of ZI agents, with humans underperforming ZI agents in terms of traded quantity less for higher market sizes. Conversely, OTC markets display the opposite comparative statics: humans' size-dependent behavior is opposite to the one of ZI agents, with humans underperforming ZI agents in terms of traded quantity more for higher market sizes. This is the first element in support of the need for a theoretical analysis of strategic incentives in OTC markets (see Section 4).

3.2 Trading price

The experimental literature on human agents suggests that the distance between trading price and equilibrium price under DA decreases as market size increases (see Chatterjee and Samuelson, 1983, and follow-up papers) and that this distance is lower than under OTC for markets of size 40 (Attanasi et al., 2016). More precisely, for $n = 40$, Attanasi et al. (2016) detect trading prices under DA closer to the equilibrium price and significantly higher than under OTC, where agents usually trade at a price lower than the equilibrium one.

Results of our simulations are reported in the left panel of Table 5, which presents the average difference between trading price and equilibrium price in markets with ZI agents for varying market sizes. Recall that by construction ZI agents cannot learn through periods. Thus, their trading price should not converge to the equilibrium one regardless of the market size. In confirmation of that, in all mechanism-size combinations of the left panel of Table 5 we detect a trading-equilibrium price average distance significantly different from zero.

Interestingly, our simulations show that while greater complexity induces a significantly greater average distance between trading and equilibrium price under DA (ANOVA, $p\text{-value}$

markets was increasing with market size (H1(OTC)) but at a slower pace than in DA markets. The rejection of H1(OTC) on the basis of traded quantity in OTC markets decreasing with market size made H1(DA-OTC) hold *a fortiori*.

¹⁶The human-ZI agents (resp., ZI-human agents) difference in the probability of observing $q < q^*$ (resp., $q > q^*$) decreases monotonically from 34 (resp., 51) percentage points for $n = 10$ markets to 13 (resp., 23) percentage points for $n = 80$ markets (t-test of the difference in percentages, $p\text{-value} < 0.01$).

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

Table 5: Average difference between trading price and equilibrium price (standard errors in brackets)

= 0.003), this does not occur under OTC, where the trading-equilibrium price average distance is invariant to the market size (ANOVA, p -value = 0.101), and never greater than under DA (it is even significantly smaller at the 10% level for $n = 20$ and at the 1% level for $n \geq 40$). Therefore, our simulations provide no evidence that under OTC a greater complexity generates a greater distortion on trading price than under DA. This finding let us predict that, with humans, the trading-equilibrium price distance in OTC should depend on market size in the same way as in DA markets, where both the theoretical and the experimental literature indicate a decreasing pattern.

This reasoning also applies when formulating the hypothesis on the DA-OTC gap as market size increases. In fact, ZI simulations provide no evidence of a worse performance of OTC compared to DA markets as market size increases – if anything, the opposite is true –, and the literature indicates that DA markets with humans converge to the equilibrium price as market size increases. With this, the DA-OTC gap in the distance between trading price and equilibrium price, experimentally detected with human subjects by Attanasi et al. (2016) for $n = 40$, should be invariant to the market size.

All this allows us to state hypothesis **H2**, which summarizes our predictions on trading prices for human agents.

H2: (DA) In DA markets the distance between trading price and equilibrium price decreases as market size increases;

(OTC) In OTC markets the distance between trading price and equilibrium price decreases as market size increases;

(DA-OTC) DA markets have trading prices closer to the equilibrium price than OTC markets, and the DA-OTC gap in the distance between trading price and equilibrium price is constant with market size.

We now test **H2** relying on the experimental results reported in the right panel of Table 5.

For DA markets, the average trading price is significantly higher than the equilibrium price for market size $n = 10$ (t-test, $p\text{-value} < 0.001$). As market size increases, the positive difference between the average trading price and the equilibrium price decreases and the former essentially coincides with the latter for the highest market size (t-test, $p\text{-value} = 0.773$). Therefore, **H2(DA) is confirmed**.

On the contrary, in OTC markets the average trading price is significantly lower than the equilibrium price independently of the size n ($p\text{-value} < 0.001$ for each n), and the absolute value of this difference is increasing in market size ($p\text{-value} < 0.01$ for each comparison 10 vs. 20, 20 vs. 40, and 40 vs. 80). With this, **H2(OTC) is rejected**, which is further evidence that in OTC markets humans perform increasingly worse as market size grows.

Finally, while for market size $n \geq 10$ there is a positive DA-OTC gap in the distance between trading price and equilibrium price, this gap becomes negative for each $n \geq 20$. With this, **H2(DA-OTC) is confirmed as for the sign of the DA-OTC gap**. However, the DA-OTC gap decreases monotonically from 2.15 for markets of size 10 to -3.28 for markets of size 80 (the difference in the trading-price distance between DA and OTC in markets with 80 agents vs. markets with 10 agents is significant: t-test, $p\text{-value} < 0.001$). With this, **H2(DA-OTC) is rejected as for the market size-dependency of the DA-OTC gap**. **ZI-human agents comparison**. We note that under OTC the human-ZI agents difference in the average distance between trading price and equilibrium price increases monotonically and significantly from -0.21 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 ZI agents in markets with 80 agents vs. markets with 10 agents is significant at the 1% level). In DA markets, we find an opposite trend, with the human-ZI agents 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 ZI agents in markets with 80 agents vs. markets with 10 agents is significant at the 1% level).

3.3 Efficiency

The existing theoretical literature on DA markets with human agents shows that efficiency increases with market size (Chatterjee and Samuelson, 1983; Satterthwaite and Williams, 1989; Gresik and Satterthwaite, 1989; Wilson, 1985). The experimental literature on human agents reports that OTC markets are less efficient than DA markets for $n = 40$ (Attanasi et al., 2016).

As in Gode and Sunder (1993) and Cason and Friedman (1996), we measure efficiency as the surplus realized from trade over the potential surplus (Market Efficiency Index). The values of this index obtained through our agent-based simulations are reported in the left panel of

Table 6, which shows that market efficiency slightly decreases as market size increases both under DA and under OTC. Therefore, DA and OTC markets with ZI agents react similarly in terms of efficiency to variations of market size, as it happened when we looked at traded quantity and trading price. With this, we have no reason against the assumption that, with humans, OTC markets respond to size as DA markets, where efficiency increases with size.

	ZI Agents		Humans	
	DA	OTC	DA	OTC
10 agents	99.1%	98.7%	93.2%	85.6%
20 agents	97.9%	97.6%	93.2%	87.2%
40 agents	97.6%	97.2%	94.9%	87.0%
80 agents	97.6%	97.1%	94.7%	82.6%

Table 6: Market efficiency index, by trading mechanism and market size

As for between-mechanism comparison, the left panel of Table 6 reports that 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). This confirms what we already know for human subjects with $n = 40$ (Attanasi et al., 2016). Moreover, given that efficiency is constant across market sizes under both mechanisms, the DA-OTC gap in efficiency is invariant to market size.

All these predictions are summarized in hypothesis **H3**.

H3: (DA) In DA markets, efficiency is increasing with market size;

(OTC) In OTC markets, efficiency is increasing with market size;

(DA-OTC) Efficiency is higher in DA markets than in OTC markets, and this DA-OTC gap is constant with market size.

We now proceed to test **H3** relying on the experimental results in the right panel of Table 6 and in Table 7. The latter reports the results of a beta regression of the efficiency index by trading mechanism over the normalized difference between traded and equilibrium quantity (from Section 3.1), the distance between average trading and equilibrium price (from Section 3.2), three dummies for market size (with size 10 as baseline), and a dummy indicating whether the sender of the accepted offer was a seller.

We find that DA markets and OTC markets behave in a very different fashion also in terms of efficiency. Specifically, the efficiency index for DA markets in the right panel of Table 6 is monotonically increasing in the market size and the dummy variables of market sizes 20, 40 and 80 in Table 7 have a significant positive effect on DA market efficiency, thereby allowing us to conclude that **H3(DA) is confirmed**.

The opposite holds for OTC market efficiency, which is significantly lower for $n = 80$ than for $n = 10$ in the right panel of Table 6, with the dummy variables of market sizes 20, 40

	DA	OTC
$(q - q^*)/q^*$	3.658 (0.000)	3.676 (0.000)
$ \bar{p} - p^* $	0.009 (0.209)	-0.003 (0.593)
20	0.310 (0.000)	-0.105 (0.013)
40	0.438 (0.000)	-0.127 (0.017)
80	0.244 (0.003)	-0.297 (0.000)
Seller Sender	0.044 (0.604)	-0.165 (0.027)
Const.	2.323 (0.000)	2.529 (0.000)

Table 7: Humans: Beta regressions of efficiency by trading mechanism over all trading periods (*p-values* in brackets)

and 80 in Table 7 having a significant negative effect. Furthermore, in OTC markets 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 second column of Table 7). This is a signal of the pressure on sellers which makes them increase their willingness to sell, ultimately leading to lower-than-equilibrium trading prices reported in Section 3.2 for each n -size OTC markets (second column of the right panel of Table 5). The Seller Sender dummy variable is instead not significant in DA markets (first column of Table 7). Therefore, **H3(OTC) is rejected**.

Finally, to compare efficiency in DA vs. OTC markets, we rely on the left column of Table 8. This reports the results of a regression of the market efficiency index over a mechanism dummy for OTC markets (with DA as baseline), three dummies for market size (with size 10 as baseline), and the interaction between the mechanism dummy and the market size dummies. This regression analysis confirms that DA markets are more efficient than OTC markets for each market size (OTC coefficient negative and significant), and that the DA-OTC gap is constant with market size for $n \leq 40$. The negative and significant coefficient of $\text{OTC} \times 80$, however, suggests that the DA-OTC gap could be increasing with size for larger market sizes. With this, we conclude that **H3(DA-OTC) is only partially confirmed**.

ZI-human agents comparison. The comparison of the left and right panels of Table 6 shows that the DA-OTC efficiency gap is significantly lower for ZI agents than for humans at any market size $n \in \{10, 20, 40, 80\}$ (t-test, highest *p-value* = 0.001). As for the size-dependent behavior, in DA markets the ZI-human agents gap in efficiency is decreasing in size from 5.88 for markets of size 10 to 2.92 for markets of size 80 (t-test, $n = 80$ vs.

	Efficiency	IM-inefficiency
OTC	-0.487 (0.000)	0.135 (0.052)
20	0.051 (0.720)	-0.040 (0.849)
40	0.522 (0.003)	-0.245 (0.368)
80	0.679 (0.002)	-1.573 (0.000)
OTC \times 20	0.131 (0.128)	0.169 (0.193)
OTC \times 40	-0.117 (0.271)	0.090 (0.562)
OTC \times 80	-0.388 (0.004)	0.682 (0.003)
Const.	2.539 (0.000)	0.344 (0.002)

Table 8: Humans: Comparison of efficiency and intra-marginal (IM) inefficiency for DA vs OTC and various market sizes (*p-values* in brackets)

$n = 10$, $p\text{-value} < 0.001$). In OTC markets, instead, the ZI-human agents gap in efficiency is decreasing monotonically with size for $n \leq 40$, from 13.14 for markets of size 10 to 10.16 for markets of size 40 (t-test, $n = 40$ vs. $n = 10$, $p\text{-value} < 0.001$), but it is increasing to 14.47 for markets of size 80 (t-test, the null hypothesis of greater or equal efficiency in $n = 80$ vs. $n = 40$ cannot be rejected, $p\text{-value} = 1.000$). We see this as further confirmation that, with human traders, OTC performance becomes problematic for large enough market sizes.

3.4 Sources of inefficiency

Following Gode and Sunder (1993) and Cason and Friedman (1996), 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*). We investigate the effects of market size on the share of IM-inefficiency.

Attanasi et al. (2016) find that, in markets with $n = 40$ human agents, inefficiency of OTC markets is a mixture of IM-inefficiency and EM-inefficiency, while inefficiency of DA markets is mostly associated with EM-inefficiency, the share of IM-inefficiency being negligible.

The left panel of Table 9 reports the results of our agent-based simulations in terms of the composition of inefficiency, i.e., the share of EM-inefficiency and the one of IM-inefficiency, by trading mechanism and market size.

Note that IM-inefficiency plays no role under DA for each market size. This is by construction

	ZI agents		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)

for DA markets of size 10, for which we chose the length of the trading period so as to allow enough time to ZI agents to reach the efficient quantity, i.e., zero IM-inefficiency (as in 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 maintained the same length of the trading period for the OTC market of the correspondent size, to check whether OTC greater complexity with respect to DA may lead to some IM-inefficiency, and how this might depend on the market size. Indeed, we detect a positive share of IM-inefficiency for OTC markets (t-test for the mean IM-inefficiency equal to 0 for each market size: the null hypothesis is rejected at the 5% level for every market size, with highest *p-value* = 0.014 for $n = 10$). However, this share of IM-inefficiency is negligible for each market size, with no detected monotonicity (it is significantly smaller than 5% regardless of the market size: the null hypothesis of a mean IM-inefficiency greater or equal to 5% is rejected, highest *p-value* < 0.001). Irrelevance of IM-inefficiency in OTC markets and its insensitivity to market size confirm that the exogenous increase of the negotiation time with market size was able to compensate for the increase in complexity of agents' interactions not only under DA but also under OTC. Therefore, moving to human agents, we assume that OTC markets behave as DA markets, i.e., IM-inefficiency decreases with market size.

Finally, as for between-mechanism comparison, given the same allowed trading time, the slight DA-OTC gap – consistently with what we already know for human agents with $n = 40$ (Attanasi et al., 2016) – seems to be independent from the market size (it increases up to $n = 40$, with *p-values* < 0.01 for both 20 vs. 10 and 40 vs. 20 positive differences, and decreases from 40 to 80, with *p-value* < 0.01).

All these predictions are summarized in hypothesis **H4**.

- H4: (DA)** In DA markets, the share of IM-inefficiency is decreasing with market size;
(OTC) In OTC markets, the share of IM-inefficiency is decreasing with market size;
(DA-OTC) The share of IM-inefficiency is lower in DA than in OTC markets, and this DA-OTC gap is constant with market size.

We test **H4** relying on the experimental results in the right panel of Table 9 and in the right column of Table 8. The latter reports the results of a regression of the share of IM-inefficiency over a mechanism dummy for OTC markets (with DA as baseline), three dummies for market size (with size 10 as baseline), and the interaction between the mechanism dummy and the market size dummies.

In DA markets, the weight of IM-inefficiency over EM-inefficiency decreases dramatically with the number of agents. 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 ($p\text{-value} = 0.519$). Therefore, **H4(DA) is confirmed**.

As for OTC markets, the weight of IM-inefficiency over EM-inefficiency does not decrease monotonically with market size. Furthermore, a difference-in-means test assesses that the difference in the shares of IM-inefficiency between size 80 and size 10 is positive and significant at the 1% level ($p\text{-value} < 0.01$), and that the difference between size 80 and size 40 is positive and significant at the 1% level ($p\text{-value} < 0.01$). Therefore, **H4(OTC) is rejected**: for OTC markets, IM-inefficiency does not disappear as market size increases. This is, once again, evidence of the strategic differences of the two trading mechanisms when used by humans.

As for the DA-OTC gap, the right column of Table 8 confirms that the share of IM-inefficiency is lower in DA than in OTC markets (OTC coefficient positive and weakly significant). Moreover, the DA-OTC gap in the share of IM-inefficiency is constant for $n \leq 40$. However, the gap seems to be increasing for larger markets (OTC \times 80 coefficient positive and significant). Hence, we conclude that **H4(DA-OTC) is only partially confirmed**.

ZI-human agents comparison. In DA markets, IM-inefficiency is null for each market size with ZI agents (second column of Table 9) and it is positive and significantly decreasing in the market size with humans (fourth column of Table 9). Hence, although the human-ZI agents gap in terms of IM-inefficiency is positive for each market size, it decreases significantly with market size under DA. In OTC markets, IM-inefficiency is small and constant for each market size with ZI agents (second column of Table 9) and it is positive but relatively stable in market size with humans (fourth column of Table 9), even significantly higher for market size 80 than for any other market size (highest $p\text{-value} < 0.001$ for the comparison with market size 10). Hence, the positive human-ZI agents gap in terms of IM-inefficiency is

not decreasing with market size under OTC. This is a final confirmation that, with human traders, OTC markets do not respond well to increases in market size.

4 Strategic interactions in OTC markets with humans

Results in Section 3 highlight a consistent pattern. While the existing theoretical literature and the predictions obtained by looking at the simulations of markets with ZI agents allow us to formulate accurate predictions of the behavior of DA markets with human agents, they do not help to describe the behavior of OTC markets with humans. In fact, our experimental results show that OTC markets respond to size differently both from DA markets with humans and from OTC markets with ZI agents. This implies that agent-based models are not the appropriate tool to predict human behavior in OTC markets as they fail to take into account how characteristics of the market, such as its size, affect the strategic behavior of agents. Therefore, in Section 4.1 we put forward a simple model that accounts for the strategic mechanisms at play in OTC markets with humans, and we show how the model explains the findings of Sections 3.1-3.4 that concern size-dependent behavior in OTC markets. Then, in Section 4.2 we further formulate a specific hypothesis (H5) which focuses on the strategic mechanisms described by our model and test it.

4.1 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 papers¹⁷ with a search model where agents with high and low valuations of an asset trade. However, in their model, 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 interaction are affected by expectations of future payoffs, which in turn are affected by the market size. We, therefore, introduce a simple bargaining model of OTC markets 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

¹⁷Duffie (2010, 2012); Ashcraft and Duffie (2007); Duffie et al. (2005, 2007); Duffie and Manso (2007); Duffie et al. (2009, 2010a,b, 2014).

equally likely.¹⁸ 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 stages.¹⁹ All the traders are present in the market at the beginning of the game. In every stage, each buyer is randomly matched to a seller (recall that there is the 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 stage in which a new random matching occurs. Proposition 1 characterizes a subgame perfect equilibrium of this game in terms of stage t offers p_t^k and acceptance strategies s_t^k , with $t = 1, 2$ and $k = \beta, \sigma$. The proof is contained in Appendix B.

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

- *Stage 1 offers are:*

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

- *Stage 1 acceptance strategies are:*

$$s_1^\beta = \begin{cases} \text{Yes} & \text{if } p_2 \leq \max\{0, \beta - 1\} \\ \text{No} & \text{otherwise;} \end{cases} \quad s_1^\sigma = \begin{cases} \text{Yes} & \text{if } p_2 \geq \min\{\sigma + 1, 4\} \\ \text{No} & \text{otherwise.} \end{cases}$$

- *Stage 2 offers are:*

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

¹⁸The theoretical analysis is the same even if we assume that buyers' valuations and sellers' costs are randomly drawn from a discrete uniform distribution 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.

¹⁹Note that what we call stage in the model is not the trading period of the experiment. The trading period in the model is the union of the two stages in which the agents interact.

- *Stage 2 acceptance strategies are:*

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

Note that in stage 1 buyers and sellers strategically reject more offers than in stage 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 delay is in line with what observed in other bargaining models with incomplete information (see, for example, Bochet and Siegenthaler, 2018; Deneckere and Liang, 2006; Ausubel et al., 2002). Our model, however, highlights a novel rationale for the existence of strategic delay, i.e., the existence of a pool of potential traders. Therefore, the effect is sensitive to their characteristics, and, importantly, to market size. Note that the strategic waiting effect decreases the number of trades in the first stage. Furthermore, offers are also higher in stage 1 than in stage 2, to compensate for lower acceptance rates. As a matter of fact, trades happen in the cases described in Table 10.

<i>Stage</i>	<i>Proposer</i>	(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 10: Feasible trades, by stages 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. Thus, only more advantageous trades happen in earlier stages, and this reduces the efficiency of the market. This explains why, in OTC markets, there is a non-negligible share of IM-inefficiency, which does not disappear as market size increases. In fact, trades between best buyers and best sellers that happen

early worsen the pool of potential traders, decrease overall traded quantity and generate inefficiency due to infra-marginal units not being traded. Importantly, this structural effect of OTC markets does not vanish with increases in market size.

To better understand the effects of size on OTC markets, we can compare the $n = 10$ market described above with an $n = 6$ market where we have 3 buyers with valuations $\beta \in \{0, 2, 4\}$ and 3 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 bad offers in the first stage; second, traders anticipate this effect and increase first-stage 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, consistently with the findings of Section 3.1 and Section 3.3, respectively.

4.2 Strategic effects in OTC markets with humans

Let us now test our strategic explanation of the behavior of human agents in OTC markets. We argued above that OTC markets with human agents are exposed to an increase in inefficiency due to the strategic behavior of traders who are less willing to accept offers in earlier interactions because they anticipate possibly higher gains from trade in future interactions. Therefore, contrary 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 interactions, 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 interactions as market size grows. We formulate hypothesis **H5** to test this strategic mechanism that underlies the model we introduced in Section 4.1.

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

Table 11 shows the average difference between trading price and equilibrium price in the first $n/5$ trades in a period, by trading mechanism and market size n . It differs from Table 5 where all trades are considered. We note that, in OTC markets, the absolute value of the average difference between trading and equilibrium price in the first $n/5$ trades is increasing in market size (second column of Table 11, p -value < 0.01 for each comparison 10 vs. 20, 20 vs. 40, and 40 vs. 80). The observation that in OTC markets the majority of the offers leading to a trade are proposed by sellers (between 55% and 61% depending on market size, significantly higher than 50% at the 1% level) helps understanding why the

average difference between trading and equilibrium price is negative for each market size. In fact, only the most advantageous asks – i.e., those being much lower than the equilibrium price – are accepted by buyers in earlier interactions. The fact that the average trading price decreases monotonically from 1 point below to 4 points below the equilibrium price as OTC market size increases from 10 to 80 agents confirms that the aforementioned strategic effect is stronger for greater market sizes.

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

Table 11: Humans: Average difference between trading price and equilibrium price for the first $n/5$ trades in a period, for markets of size $n \in \{10, 20, 40, 80\}$ (standard errors in brackets)

As a control, the first column of Table 11 shows that the same does not occur for DA markets. In fact, the average difference between trading and equilibrium price for the first $n/5$ trades 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 trade due to sellers’ asks is not significantly higher than the fraction due to buyers’ bids (t-test of the difference in the two average fractions of offers, p -value = 0.447). With this, we conclude that **H5 is confirmed**.

5 Conclusions

The sequential bargaining approach analyzed in this paper is an over-the-counter (OTC) market as it “does not use a centralized trading mechanism, such as an auction, specialist, or limit-order book, to aggregate bids and offers and to allocate trades.” (Duffie, 2012, p. 20).²⁰ Thus, price transparency is lower than in exchange markets. In our setting, in particular, only post-trade price transparency is enforced, but agents are in the dark about the full history of asks and bids. With this, our main focus is on behavior of OTC markets of varying size both when the interaction occurs among humans and when the interaction occurs among zero-intelligence (ZI) agents. To the best of our knowledge, this is the first

²⁰For equivalent definitions of OTC markets and a discussion of the economic relevance of these trading mechanisms, see, e.g., Duffie et al. (2005), Duffie et al. (2007), Hendershott and Madhavan (2015).

theory-driven experimental study of OTC markets that compares such decentralized market mechanisms with centralized trading institutions (DA), controlling for human vs. ZI agents performance under different market sizes.

Exploring electronic OTC mechanisms with public information about trading prices is worth doing. 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) (see Bessembinder and Maxwell, 2008; Edwards et al., 2007; Goldstein et al., 2006; Green et al., 2007a,b, 2010, for empirical analyses).²¹ The increase in OTC financial transactions “would not have been possible without the dramatic advances in information and computer technologies that have occurred” from 1980 to 2000 (Schinasi et al., 2000, p. 1). The increasingly common use of OTC trading mechanisms in today’s digital world thus justifies the focus of our study.

Our paper highlights how inefficiencies in OTC markets originate not only from the higher complexity of the market structure but also from the strategic interaction between agents, in a twofold manner.

The first contribution is experimental. Relying on a large experimental dataset of 6,400 undergraduate students in Economics and Management, we show that while DA markets with humans behave according to the predictions derived from the theoretical literature and the agent-based simulations with ZI agents, OTC markets do not. The simulations with ZI agents show that these results are not merely due to OTC market rules given that 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. Rather, the result is due to the fact that human agents’ strategic response differs across market structures.

The second contribution is theoretical: in the literature, OTC markets have often been 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

²¹However, differently from the OTC market of our experiment, in these OTC markets, agents usually trade multiple units of the asset; and the traded quantity is reported only if lower than a stipulated threshold. Other OTC markets include credit default swaps markets, in which traders can select from a menu of bids and asks, and therefore have a different information set than in our experimental market (Feldhütter, 2012).

bargaining model that explicitly describes agents' strategic interaction in OTC markets.²² In fact, not only learning but also strategic interaction distinguishes human agents from ZI agents' behavior. As for the latter, humans' bargaining strategies in every period are affected by expectations of future payoffs, which in turn are affected by market size. Relying on this intuition, the model explains our main experimental findings – lower trading quantity, higher price dispersion and lower market efficiency when OTC market size increases – which were not in line with the behavior of markets with ZI agents. Furthermore, with human traders, the DA-OTC efficiency gap increases as 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 interactions decrease and earlier offers increase, but the second effect is not always enough to compensate the decrease in acceptance rates. It may be interesting to explore theoretically and experimentally whether our results still hold when goods are heterogeneous along the lines of Bos (2019).

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).²³ However, we acknowledge that monetary incentives are important in market experiments. On the other side, we stress the point that our study is comparative: behavior in OTC markets is analyzed in contrast to behavior in DA markets and behavior of human agents in OTC markets is analyzed in contrast to behavior of ZI agents in OTC markets. 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).

²²Indeed, negotiations in OTC markets may be interpreted as a combination of bilateral seller-buyer negotiations, hence several of the theoretical features of bilateral and multilateral non-cooperative bargaining models apply (see Gomes, 2020; Konrad and Thum, 2020, and references therein).

²³For 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.

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 regardless of the market size. 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 trades 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.

Finally, both our theoretical and experimental results show that large OTC markets with human agents have inefficiency problems. This is a serious concern, as typically real-life applications of OTC markets have a size substantially larger than those of our experimental markets (e.g., OTC derivatives markets totaled approximately US\$601 trillion in 2010). Therefore it would be interesting to investigate whether changes in the design of OTC markets may soften this issue. One possible suggestion is the introduction of some form of assortative matching, which would be easily implemented in experiments, but more hardly so in real markets. An alternative route could be the introduction of an endogenous matching procedure in the spirit of Kim (2012), which could induce assortative matching endogenously. Intuitively, these matching procedures should help increase the efficiency of large OTC markets in a twofold manner: first, they should save agents' searching time, making it more likely that an agent meets a suitable counterpart. Second, they should prevent early trade between the best agents, i.e., buyers with the highest valuations and sellers with the lowest costs. Due to the agents' strategic behavior, these are the only trades that can happen early in the period, but when they happen they significantly worsen the pool of agents, making further trade less likely. Whether these theoretical intuitions will be confirmed experimentally is a relevant question for future research.

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 Lefebvre. Controlling monopoly power in a classroom double-auction market experiment. Working Paper No. 2019-08, Bureau d’Economie Théorique et Appliquée, UDS, Strasbourg, 2019.
- Lawrence M Ausubel, Peter Cramton, Raymond J Deneckere, et al. Bargaining with incomplete information. *Handbook of Game Theory*, 3:1897–1945, 2002.
- Nicholas Bardsley, Robin Cubitt, Graham Loomes, Peter Moffatt, Chris Starmer, and Robert Sugden. *Experimental Economics: Rethinking the Rules*. Princeton, NJ: Princeton University Press, 2009.
- Hendrik Bessembinder and William Maxwell. Markets: Transparency and the corporate bond market. *Journal of Economic Perspectives*, 22(2):217–234, 2008.
- Max R. Blouin and Roberto Serrano. A decentralized market with common values uncertainty: Non-steady states. *Review of Economic Studies*, 68:323–346, 2001.
- Olivier Bochet and Simon Siegenthaler. Better later than never? an experiment on bargaining under adverse selection. *International Economic Review*, 59(2):947–971, 2018.
- Olivier Bos. Charitable asymmetric bidders. *Journal of Public Economic Theory*, 22(2):320–337, 2019.
- Paul J. Brewer. Zero-intelligence robots and the double auction market: A graphical tour. In Charles R. Plott and Vernon L. Smith, editors, *Handbook of Experimental Economics Results*, volume 1, pages 31–45. Elsevier, 2008.
- 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. *Journal of Political Economy*, 56(2):95–108, 1948.
- Kalyan Chatterjee and William Samuelson. Bargaining under complete information. *Operations Research*, 93:835–851, 1983.
- Joaquín Coleff and Daniel Garcia. Information provision in procurement auctions. *Journal of Public Economic Theory*, 19(2):426–444, 2017.
- James C Cox, Mark Rider, and Astha Sen. Tax incidence: Do institutions matter? an experimental study. *Public Finance Review*, 46(6):899–925, 2018.
- Martin W Cripps and Jeroen M Swinkels. Efficiency of large double auctions. *Econometrica*, 74(1):47–92, 2006.
- Raymond Deneckere and Meng-Yu Liang. Bargaining with interdependent values. *Econometrica*, 74(5):1309–1364, 2006.
- Darrell Duffie. Asset Price Dynamics with Slow-Moving Capital. *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.
- Amy K. Edwards, Lawrence E. Harris, and Michael S. Piwowar. Corporate Bond Market Transaction Costs and Transparency. *Journal of Finance*, 62(3):1421–1451, 2007.
- Peter Feldhütter. The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures. *Review of Financial Studies*, 25(4):1155–1206, 2012.
- 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 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? *Quarterly Journal of Economics*, 112(2):603–630, 1997.
- Dhananjay 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. Lower bounds for efficiency of surplus extraction in double auctions. In *The Double Auction Market*, pages 199–220. Routledge, 2018.
- Michael A. Goldstein, Edith S. Hotchkiss, and Erik R. Sirri. Transparency and Liquidity: A Controlled Experiment on Corporate Bonds. *Review of Financial Studies*, 20(2):235–273, 2006.
- Armando Gomes. Valuations and dynamics of negotiations. *Journal of Public Economic Theory*, 22(1):245–273, 2020.
- Richard C. Green, Burton Hollifield, and Norman Schürhoff. Dealer intermediation and price behavior in the aftermarket for new bond issues. *Journal of Financial Economics*, 86(3): 643 – 682, 2007a.

- Richard C. Green, Burton Hollifield, and Norman Schürhoff. Financial intermediation and the costs of trading in an opaque market. *Review of Financial Studies*, 20(2):275–314, 2007b.
- Richard C. Green, Dan Li, and Norman Schürhoff. Price Discovery in Illiquid Markets: Do Financial Asset Prices Rise Faster Than They Fall? *Journal of Finance*, 65(5):1669–1702, 2010.
- 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.
- Terrence Hendershott and Ananth Madhavan. Click or call? auction versus search in the over-the-counter market. *Journal of Finance*, 70(1):419–447, 2015.
- Charles A. Holt. Classroom Games: Trading in a Pit Market. *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.
- Kyungmin Kim. Endogenous market segmentation for lemons. *RAND Journal of Economics*, 43(3):562–576, 2012.
- Kai A Konrad and Marcel Thum. Equilibrium opacity in ultimatum-offer bargaining. *Journal of Public Economic Theory*, 22(5):1515–1529, 2020.
- Stephan Laueremann. Asymmetric information in bilateral trade and in markets: An inversion result. *Journal of Economic Theory*, 147:1969–1997, 2012.
- 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.
- Diego Moreno and John Wooders. Prices, delay, and the dynamics of trade. *Journal of Public Economic Theory*, 104:245–273, 2002.
- Andrea Morone, Francesco Nemore, and Simone Nuzzo. Experimental evidence on tax salience and tax incidence. *Journal of Public Economic Theory*, 20(4):582–612, 2018.

- 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.
- Bradley J Ruffle. Competitive equilibrium and classroom pit markets. *Journal of Economic Education*, 34(2):123–137, 2003.
- 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. *Review of Economic Studies*, 56(4):477–498, 1989.
- Garry J. Schinasi, R. Sean Craig, Burkhard Drees, and Charles Kramer. Modern banking and otc derivatives markets the transformation of global finance and its implications for systemic risk. *IMF Occasional Papers*, 103, 2000.
- Vernon L. Smith. An Experimental Study of Competitive Market Behavior. *Journal of Political Economy*, 70(2):111–137, 1962.
- Donald A Wells. Laboratory experiments for undergraduate instruction in economics. *Journal of Economic Education*, 22(3):293–300, 1991.
- Robert Wilson. Incentive efficiency of double auctions. *Econometrica*, 53(5):1101–1115, 1985.
- Robert B Wilson. On equilibria of bid-ask markets. In *Arrow and the Ascent of Modern Economic Theory*, pages 375–414. Springer, 1987.
- John Wooders. Matching and bargaining models of markets: approximating small markets by large markets. *Economic Theory*, 11:215–224, 1998.

Appendix

Appendix A: Details of simulations

A.1 Simulation of Double Auction markets with ZI agents

We set a bid-ask improvement rule as in Brewer (2008, pp. 32-34). The buyer with the current best bid and the seller with the current best ask close the trade. 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.²⁴ 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 the market clears. If the bid is lower than the current ask, it remains as the best bid, and the agent becomes inactive.

²⁴The 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.

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:

- 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 the 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.

A.2 Simulation of Over-the-Counter markets 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 2008, 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 2008, 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 offer 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.

Appendix B: Theoretical Appendix

B.1 Proof of Proposition 1

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

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

Consider now the offer of the second stage.

Second stage: 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/she can only trade if he/she meets the seller with cost 0 and he/she 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/she can trade only with sellers of type $\sigma \in \{0, 1\}$. If he/she proposes $p_2 = 1$ he/she makes zero profit, hence he/she proposes $p_2(\beta = 1) = 0$ and makes expected profit $\mathbb{P}_2[\sigma = 0]$.

- A buyer with valuation $\beta = 2$, knows he/she can trade with sellers of type $\sigma \in \{0, 1, 2\}$. If he/she proposes $p_2 = 2$ he/she makes zero profit. If he/she proposes $p_2 = 1$ he/she 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/she proposes $p_2 = 0$ he/she only trades with sellers of type $\sigma = 0$ and makes profit 2 when doing so, hence his/her expected profit is $2 \cdot \mathbb{P}[\sigma = 0]$. Therefore, he/she 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/she can trade with sellers of type $\sigma \in \{0, 1, 2, 3\}$. If he/she offers $p_2 = 3$ he/she makes zero profits; if he/she offers $p_2 = 2$ he/she makes an expected profit of $\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1] + \mathbb{P}_2[\sigma = 2]$; if he/she offers $p_2 = 1$ he/she makes an expected profit of $2(\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1])$; finally, if he/she offers a price $p_2 = 0$ he/she makes an expected profit of $3 \cdot \mathbb{P}_2[\sigma = 0]$. Therefore, he/she 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/she can trade with all sellers. If he/she offers $p_2 = 4$ he/she makes zero profits; if he/she offers $p_2 = 3$ he/she 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/she offers $p_2 = 2$ he/she makes an expected profit of $2(\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1] + \mathbb{P}_2[\sigma = 2])$; if he/she offers $p_2 = 1$ he/she makes an expected profit of $3(\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1])$; finally, if he/she offers a price $p_2 = 0$ he/she makes an expected profit of $4 \cdot \mathbb{P}_2[\sigma = 0]$. Therefore, he/she will offer

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

Second stage: sellers' offers. 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

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

First stage: acceptance strategies. Note that the second-stage 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 stage 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 $2(\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 1])$, 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 stage. As a consequence, choosing $p_2 = 2$ leads to the (weakly) highest payoff of 2 ($\sum_{k=0}^2 \mathbb{P}[\sigma = k]$), 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 stage 2. As for the expected payoff from receiving an offer, the maximum expected payoff from stage 2 is $\frac{4}{5}$, which happens when all the sellers of type $\sigma \in \{0, 1, 2\}$ are still in the market in stage 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 stage, the acceptance strategies are:

$$s_1^\beta = \begin{cases} \text{Yes} & \text{if } p_2 \leq \max\{0, \beta - 1\} \\ \text{No} & \text{otherwise;} \end{cases} \quad s_1^\sigma = \begin{cases} \text{Yes} & \text{if } p_2 \geq \min\{\sigma + 1, 4\} \\ \text{No} & \text{otherwise.} \end{cases}$$

Consider now the offer of the first stage.

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

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

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

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

■

B.2 The case of $n = 6$ and the effect of market size

Section 4.1 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).²⁵

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 for the case $n = 10$ (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 bargaining game.*

- Stage t offers, for $t = 1, 2$, are:

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

- Stage 1 acceptance strategies are:

$$s_1^\beta = \begin{cases} \text{Yes} & \text{if } p_2 \leq \max\{0, \beta - 2\} \\ \text{No} & \text{otherwise,} \end{cases} \quad s_1^\sigma = \begin{cases} \text{Yes} & \text{if } p_2 \geq \min\{\sigma + 2, 4\} \\ \text{No} & \text{otherwise.} \end{cases}$$

²⁵Note that the relation between market size and efficiency crucially depends on the assumptions on how market size changes.

- *Stage 2 acceptance strategies are:*

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

Proof:

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

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

Consider now the offer of the second stage.

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

- A buyer with valuation $\beta = 0$ knows that he/she can only trade if he/she meets the seller with cost 0 and he/she 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/she can trade only with sellers of type $\sigma \in \{0, 2\}$. If he/she proposes $p_2 = 2$ he/she makes zero profit, hence he/she 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/she can trade with all sellers. If he/she offers $p_2 = 4$ he/she makes zero profits; if he/she offers $p_2 = 2$ he/she makes an expected profit of $2 (\mathbb{P}_2[\sigma = 0] + \mathbb{P}_2[\sigma = 2])$; finally, if he/she offers a price $p_2 = 0$ he/she makes an expected profit of $4 \cdot \mathbb{P}_2[\sigma = 0]$. Therefore, he/she 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 stage: sellers' offers. 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 stage: acceptance strategies. Note that the second-stage 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 stage 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 stage. 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

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

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 stage, the acceptance strategies are:

$$s_1^\beta = \begin{cases} \text{Yes} & \text{if } p_2 \leq \max\{0, v(\beta) - 2\} \\ \text{No} & \text{otherwise;} \end{cases} \quad s_1^\sigma = \begin{cases} \text{Yes} & \text{if } p_2 \geq \min\{c(\sigma) + 2, 4\} \\ \text{No} & \text{otherwise.} \end{cases}$$

Consider now the offer of the first stage.

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

- A buyer with low valuation, $\beta = 0$, knows that he/she can never trade in the first stage. Therefore he/she proposes $p_1(\beta = 0) = 0$, as any other price may induce a negative payoff.
- A buyer with intermediate valuation, $\beta = 2$, knows he/she can trade if he/she meets a seller with low cost. However, his/her only opportunity to trade with a seller of type $\sigma = 0$ is to offer $p_1 = 2$. In this case, he/she would forego the possible positive payoffs of stage 2, therefore $p_1(\beta = 2) = 0$.
- A buyer with high valuation, $\beta = 4$, knows he/she can trade with any potential seller. However, in order to trade with sellers of type $\sigma \in \{2, 4\}$ he/she has to offer $p_1 = 4$

which gives zero payoff. Given that his/her expected payoff from stage 2 is positive and smaller than 2, his/her optimal choice is to offer $p_1 = 2$ which gives him/her a payoff of 2 in case the offer is accepted (which happens with probability $\mathbb{P}_1[\sigma = 0] = \frac{1}{3}$). Therefore $p_1(\beta = 4) = 2$.

First stage: sellers' offers. We can symmetrically derive the sellers' optimal offers in the first stage, 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-stage 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 stage 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 stage when the market size increases. Second, traders anticipate this effect and modify first-stage 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).²⁶

²⁶The expected performance has been computed by means of a Python code which is available upon request from the authors.

	$n = 6$	$n = 10$
$q < q^*$	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 theoretical OTC market, by market size

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 stage 1 – is stronger than the second one – i.e., buyers’ (resp., sellers’) offers being weakly higher (resp., lower) in stage 1.