



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

European Economic Review 49 (2005) 1827–1853

EUROPEAN
ECONOMIC
REVIEW

www.elsevier.com/locate/econbase

Endogenizing market institutions: An experimental approach

Georg Kirchsteiger^a, Muriel Niederle^b, Jan Potters^{c,*}

^a*Department of Economics, University of Maastricht, P.O. Box 616, 6200 MD Maastricht, the Netherlands*

^b*Department of Economics, Stanford University, Stanford, CA 94305-6072, USA*

^c*Tilburg University, CentER/ Department of Economics, P.O. Box 90153, 5000 LE Tilburg, the Netherlands*

Received 12 February 2003; accepted 25 May 2004

Available online 23 August 2004

Abstract

We study experimental two-sided markets in which the information structure is endogenous. When submitting an offer, a trader decides which other traders will be informed about the offer. This setup allows both a decentralized bargaining market (Chamberlin, *J. Polit. Econ.* 56 (1948) 95), and a double auction market (Smith *J. Polit. Econ.* 70 (1962) 111) as special cases. The results show that offers are typically directed to all traders of the *other* side of the market, but to none of the traders of the *same* side of the market. Even though traders receive much less information, the resulting market institution leads to the same outcomes in terms of prices and efficiency as a double auction market. In two additional treatments we examine the robustness of these results. First, it is found that the market institution adapts predictably, but not necessarily efficiently, to the imposition of transaction costs. Second, we find that the preference of sellers to conceal offers from competitors is strict. At the same time, sellers benefit collectively when they reveal offers to each other.

© 2004 Elsevier B.V. All rights reserved.

JEL classification: C9; D4; L1

Keywords: Market institution; Matching and information structures; Efficiency

*Corresponding author. Tel.: +31-13-4668204; fax: 31-13-4883042.

E-mail address: j.j.m.potters@uvt.nl (J. Potters).

1. Introduction

One of the main lessons of experimental economics is that market institutions matter. The rules that govern trade may exert a profound effect on prices, efficiency and the speed of convergence (Plott, 1982; Holt, 1995). In view of this insight important questions are: Which market institution is likely to emerge endogenously? Can we be confident that an efficient institution will emerge?

The function of market institutions is to match potential trading partners and to provide them with the information needed to form contracts. Therefore, differences in the information structure and the matching procedure can have a substantial impact on the performance of markets. This is nicely illustrated by two seminal experimental studies. Chamberlin (1948) studied a decentralized bargaining market (DBM) in which traders contacted each other one by one, exchanging offers on a bilateral basis, with no other trader being informed of these offers or of realized trades. This market was found to generate inefficient outcomes inconsistent with the competitive equilibrium. Smith (1962) studied a double auction market (DAM) in which traders submitted offers on a multilateral basis, with all traders being informed of these offers as well as of realized trades. This market was found to converge quickly to the competitive equilibrium.

The present paper studies experimental markets in which the structure of the information flows is endogenous. When a trader submits an offer, she decides whom to inform of the offer. She can decide to inform any subset of traders from both the own side of the market and the other side of the market. If a trader is informed of an offer from the other side of the market, she can form a contract by accepting the offer. Only traders that are informed of the originating offer are informed of such a contract. Note that if each trader always chooses to inform every other trader of her offer, the information and matching structure is identical to that of a DAM. If, at the other extreme, each trader always chooses to inform only one other trader of her offer, the information and matching structure is in essence that of a DBM. Furthermore, not only DAM and DBM are special cases of our setup. For example, also an offer auction (only sellers make offers, all traders are informed), a bid auction (only buyers make offers, all traders are informed), a posted offer auction (sellers post take-it-or-leave offers, all traders are informed), and a posted bid auction (buyers make take-it-or-leave offers, all traders are informed) are institutions that could emerge in our design.¹

We first investigate the endogenously emerging institution in a ‘pure’ environment without exogenous frictions. The results indicate that offers are made by both sides of the markets, and that offers from both buyers and sellers are typically revealed to all potential trading partners but are concealed from rivals. Hence, traders receive less information than in the DAM and more than in the DBM. Nonetheless, the

¹An interesting related experimental study is Campbell et al. (1991), where traders could choose to trade at a double auction or by means of bilateral bargaining. Campbell et al. (1991) however, did not allow for intermediate institutions, which may be important as we will see later.

market outcomes share the properties of the DAM in terms of efficiency and price convergence.

To further investigate these results we employ two additional treatments. First, we impose costs on the buyers for sending offers. The results indicate that sellers now make most of the offers. These offers are again transmitted to all buyers but concealed from other sellers. The emerging market is less efficient than the DAM. We conclude that transactions costs can indeed shape the emerging market institution, but that there is no guarantee that this institution will lead to efficient outcomes.

Second, we further examine the preference to conceal offers from competitors. We use a treatment in which every seller receives a small subsidy if he chooses to reveal an offer to his rivals. The results indicate that a large fraction of the sellers still chooses to conceal offers from rival sellers. Hence, the preference to conceal offers from competitors is strict rather than weak. At the same time, the results indicate that in those instances in which the sellers choose to reveal offers to each other prices tend to be higher.

The paper is organized as follows. The next section describes the features of the experimental design common to all treatments. In Section 3 the endogenous market without frictions is investigated. Section 4 discusses the treatment with transaction costs imposed on buyers. Section 5 investigates the effect of a price communication subsidy for sellers. The concluding discussion can be found in Section 6.

2. Experimental design

2.1. Market structure

Each experimental market session consisted of a sequence of one practice round and 18 trading rounds. Each trading round lasted 3 minutes. Markets were inhabited by 12 traders, 6 buyers and 6 sellers. Traders retained their roles throughout the session. In a trading round, each individual trader could trade at most one unit of a homogenous good at no trading costs. The private value of a trader (cost value or redemption value) changed from round to round. The set of values, however, remained the same and this was common knowledge. Hence, the market environment (i.e., the induced aggregate demand and supply function) was constant across the practice round and the 18 rounds of the experiment. The cost and redemption values were given in points, and 1 point exchanged for 0.3 Dutch guilders (Hfl).²

The demand and supply functions induced by these values are illustrated in Fig. 1. As can be seen, the competitive equilibrium was at a quantity of 4 units and a price in

²1 Hfl exchanged for about € 0.45.

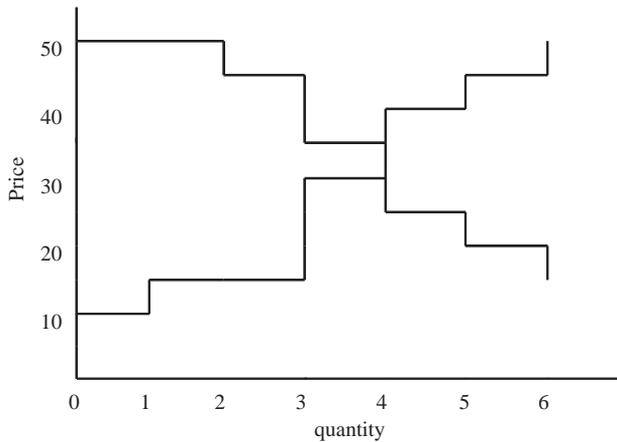


Fig. 1. The induced demand and supply schedule.

the interval 30, 35. Hence, for both buyers and sellers there were four infra- and two extra-marginal traders. The induced demand and supply schedule were, of course, to a large extent arbitrary. Our main consideration was to strike a balance between the probability of an extra-marginal trade occurring (a high probability requiring elastic schedules), and the efficiency loss associated with an extra-marginal trade (a large efficiency losses requiring inelastic schedules).

Besides values, traders were also assigned ID-letters: A, B, C, D, E, and F for buyers, and U, V, W, X, Y, and Z for sellers. These ID-letters were randomly (re)assigned to the traders at the beginning of each round. The assignment of ID-letters was neither related to the assignment of values nor to the “real” identity of the subjects. The subjects were informed about this.

2.2. Information display

The trading process was handled by means of networked computers. At any time a trader’s computer screen displayed the following information: the round number, the time left for trading, a trader’s role (buyer or seller) and ID-letter, the cost or resale value, and a trader’s total profits up to that round. In treatment with transaction costs, the costs a buyer had already incurred during the round at hand was also displayed on the screen.

In the middle of the display were the lists of ask- and bid prices, one above the other. These lists only contained those offers that the trader was informed about. Both ask- and bid prices were ordered from high to low, and for each offer also the ID-letter of the sender was indicated.

Finally, at the bottom of the screen, there was a row with the prices of those accepted offers that the trader had been informed about. Only prices were revealed but not the IDs of the transacting parties.

2.3. Procedure

For each treatment,³ four independent experimental sessions were run. Students at Tilburg University were recruited as subjects through announcement in the university bulletin and in classes. Participants were solicited for a two hours decision-making experiment which would earn them money. Fifteen subjects were registered for each experimental session to allow for no-shows. In session DAM1, however, only 10 subjects showed up. This session was run with 10 traders, using the design of Fig. 1 with one buyer (redemption value 50) and one seller (cost value 15) excluded, leaving the range of equilibrium prices unchanged.

Upon arrival in the lab, subjects drew an envelope with a seat number. If more than 12 subjects showed up, one to three empty envelopes were added to the stack of seat numbers. The subjects drawing an empty envelope received 10 Hfl for showing up and left the room. Once the remaining subjects were seated, the instructions for the experiment were distributed and read aloud by the experimenter. Then the subjects were given some minutes to study the instructions at their own pace, and to privately ask questions. After the practice round, the 18 rounds which determined subjects' earnings were run. After round 18 the subjects privately received their total earnings and left the room. Sessions lasted about one and a half hours. Earnings were on average about Hfl 49, and ranged between Hfl 21 and 65.

3. The market without frictions

In this section, we investigate the market institution that emerges endogenously when the traders decide about the direction of their offers. To assess the properties of this institution we compare it to investigate three different treatments.

3.1. Trading institutions

Treatment 1: directed bid–ask market (DBAM)

Buyers could try to buy by making bids, and sellers could try to sell by making asks. When making an offer (an ask or a bid), a trader had to enter a price at which she was prepared to trade. She also had to enter the IDs of those traders she wanted to inform about the offer. We call this treatment DBAM. This expresses that bids and asks were 'directed' (to one trader, to all traders, or to any number of traders in between) and that DAM and DBM are special cases. A trader was forced to enter the ID of at least one trader from the other market side, i.e. the ID of at least one potential trading partner. This constraint was imposed to prevent traders from sending 'fake' offers to only their own market side, i.e. to only their competitors. Furthermore, offers that could lead to negative profits were not permitted.

Offers could be adjusted at any moment by simply submitting a new offer. The new offer did not have to be an offer that is better for the other market

³The treatments will be described below.

side.⁴ With the new offer the old one became invalid. Hence, each trader had at most one outstanding offer. Since everyone could trade at most one unit, a trader who had already made a transaction could no longer make offers.

As long as a trader had not traded in that round, she was allowed to accept any offer of a potential trading partner about which she was informed. Again we enforced the restriction that an offer could only be accepted if it led to a non-negative profit. We did, however, not enforce the rule that a trader always had to accept the best price offered. When an offer was accepted, it was withdrawn from the market, and those traders who had been informed about the offer were also informed that a transaction had occurred at that price.

Treatment 2: double auction market (DAM)

In the DAM treatment, traders did not have to enter any trader ID-letters when submitting an offer. All offers were automatically sent to all other traders, i.e. to all potential trading partners as well as to all competitors. All traders were also informed when and at what price a trade occurred. This treatment boiled down to the standard continuous DAM with the exception that trader IDs were added and that traders were restricted neither to make an improving offer nor to accept the best available offer.

Treatment 3: decentralized bargaining market (DBM)

In the DBM treatment, traders had to enter one and only one ID-letter of a potential trading partner when they entered an offer. Furthermore, when a transaction occurred, only the two parties involved were informed. As in the other treatments, traders were restricted neither to make an improving offer nor to accept the best available offer.

An overview of the treatments investigated in this paper is given by [Table 1](#). The last two treatments (DBAM-TC and DBAM-PC) will be explained in detail in Sections 4 and 5. The middle column of [Table 1](#) describes whether a treatment is characterized by an exogenous or an endogenous market institution and what its main features are. The third column describes the type of friction used in each treatment.

3.2. Expected results

To describe the (expected) results it is useful to introduce the following two variables. D_{client} refers to the ‘dissemination of an offer among clients’, and is defined as the fraction of potential trading partners informed about an offer. D_{comp} refers to the ‘dissemination of an offer among competitors’ and is defined as the fraction of

⁴We did not impose an improvement rule as traders are not necessarily informed of all other outstanding offers. To keep the rules and experimental instructions identical across treatments we did not impose an offer improvement rule in any of our treatments. Note that each trader can effectively retract an offer by entering a new offer with an unacceptable price (eg., 0 for buyers or 100 for sellers).

Table 1
The treatments

Treatment	Type of institution	Type of friction
Double Auction Market (DAM)	Exogenous: all buyers and sellers are informed of all bids and asks	No frictions
Decentralized Bargaining Market (DBM)	Exogenous: only one buyer (seller) is informed of an ask (bid)	No frictions
Directed Bid–Ask Market (DBAM)	Endogenous: buyer (seller) decides who is informed of her bid (ask)	No frictions
Directed Bid–Ask Market with Transaction Costs (DBAM-TC)	Endogenous: buyer (seller) decides who is informed of her bid (ask)	Transaction costs for buyers no transaction costs for sellers
Directed Bid–Ask Market with Price Communication (DBAM-PC)	Endogenous: buyer (seller) decides who is informed of her bid (ask)	Transaction costs for buyers price communication subsidy for sellers

competitors (i.e., other traders on the same side of the market) informed about an offer. These two variables capture the defining features of the different trading mechanisms. For example, in a DAM we have $D_{\text{client}} = 1$ and $D_{\text{comp}} = 1$ for all offers. All traders are informed about all asks and bids. In a DBM we have $D_{\text{client}} = 1/6$ and $D_{\text{comp}} = 0$ for all offers. An offer is sent to one trader on the other market side, and no other trader is informed about it. In the endogenous market treatment the two variables were endogenous and determined by the traders themselves.

For reasons of symmetry, we did not expect to find any difference between the demand and the supply side of the market in the DBAM treatment. That is, we expected that the numbers of asks and bids were the same, and that $D_{\text{client}}(\text{asks}) = D_{\text{client}}(\text{bids})$ and $D_{\text{comp}}(\text{ask}) = D_{\text{comp}}(\text{bids})$. Moreover, there did not seem to be a compelling reason for traders to send an offer to only a subset of the potential trading partners. Sending an offer to all potential trading partners could only increase the probability of acceptance. Furthermore, offers that were not yet accepted could always be adjusted in any direction. Since each trader could trade at most one unit in our design, arguments for price discrimination do not apply. Hence, we expected subjects to transmit their offers to every potential trading partner: $D_{\text{client}} = 1$. However, it was not a priori clear what to expect regarding the information to competitors. On the one hand, traders may want to conceal an offer from competitors ($D_{\text{comp}} = 0$) in order to prevent or delay being overbid (respectively, undercut) by competitors. The resulting institution would then be characterized by offers which apply to all potential trading partners ($D_{\text{client}} = 1$) but

which are kept secret from competitors ($D_{\text{comp}} = 0$). Such an institution could be referred to as a ‘sealed offer double auction’, expressing the fact that offers are sent by both sides of the market like in the double auction, but that offers are concealed from rivals. On the other hand, traders may choose to inform rivals about their offers in an attempt at collusion. When competitors know each others’ offers, secret rebates are excluded. The incentive to outbid or undercut may be reduced when rivals can perfectly observe and follow them. After all, “the policing of a price agreement involves an audit of the transaction prices” (Stigler, 1964, p. 47). If competitors decide to inform each other about offers ($D_{\text{comp}} = 1$) then, together with $D_{\text{client}} = 1$, the resulting institution would be equivalent to a DAM.

What did we expect regarding prices and efficiency? As in numerous other DAM experiments, we expected prices and quantities in our DAM treatment to converge quickly to the market clearing level with outcomes close to full efficiency.⁵ For the DBM treatment, on the other hand, we expected a slower convergence of prices and quantities and lower efficiency levels.

For the DBAM treatment, we expected the outcomes to be conditional on the results regarding D_{comp} . If the directed bid–ask market would lead to an institution that is equivalent to a double auction market, we expected prices, efficiency, and convergence to be in line with the double auction market treatment. If, on the other hand, a sealed offer double auction market would emerge then it is more difficult to predict its outcomes. Such a market has never been investigated before. Nevertheless, some clear hints are provided by several bounded rationality models that are developed to explain convergence of prices and quantities in DAMs (Easley and Ledyard, 1993; Friedman, 1991; Gode and Sunder, 1997; Gjerstad and Dickhaut, 1995). In these models traders respond adaptively to observations of offers of competitors and potential trading partners and realized transactions prices. Furthermore, in these models the speed of convergence depends positively on (the number of) these observations. In a sealed offer, DAM traders would typically observe only half of the offers and realized prices, whereas these data are public information in the DAM. Hence, a reasonable expectation would be that convergence of prices and quantities to the market clearing level is slower in the sealed offer double auction market than in the DAM, but still faster than in the DBM.

3.3. Results

First we examine which market institution actually emerged in the directed bid–ask market with no frictions (Result 1). Then we compare prices (Result 2) and efficiency levels (Result 3) across the three different treatments (DBAM, DAM and DBM).

⁵Our DAM differs however from most other implementations of a double auction market, since we did not enforce an offer improvement rule. While this rule is important for convergence of double auction trades using simulated agents (see Gode and Sunder, 1997), we did expect the results of the double auction institution to be robust to that change.

Table 2
Dissemination of offers among clients (D_{client}) and competitors (D_{comp})

Session	Average	Median	% offers with	Average	Median	% offers with
	D_{client}	D_{client}	$D_{\text{client}} = 1$	D_{comp}	D_{comp}	$D_{\text{comp}} = 0$
DBAM1	0.69	1	55	0.05	0	93
DBAM2	0.92	1	87	0.09	0	89
DBAM3	0.94	1	88	0.10	0	87
DBAM4	0.87	1	82	0.24	0	74
All sessions	0.85	1	77	0.12	0	85

Result 1. Market institution. *The market institution emerging in the DBAM treatment is best described as a sealed offer double auction market. Offers were typically sent to every potential trading partner but to no competitor.*

In order to present the evidence for Result 1 we calculated for each offer the value of D_{client} , that is, the number of potential trading partners to whom the offer was sent, divided by the maximum number to whom it could have been sent (i.e. divided by 6).⁶ Notice that D_{client} ranged between 1/6 and 1 since each offer had to be sent to at least one potential trading partner. Similarly, we computed D_{comp} for each offer, that is, the number of competitors who were informed about the offer divided by the maximum number who could have been informed (i.e. divided by 5). Of course, D_{comp} ranged between 0 and 1.

Table 2 presents mean and median values for both D_{client} and D_{comp} for each of the four DBAM sessions separately and averaged over all rounds. Also the percentage of offers with $D_{\text{client}} = 1$ and $D_{\text{comp}} = 0$ is indicated. The table shows that an offer was sent to 85% of the potential trading partners on average, with 100% being the median value. Furthermore, 77% of all offers were sent to every potential trading partner ($D_{\text{client}} = 1$). A look at the session data indicates that this pattern was representative also for the individual sessions. Even in DBAM1, where average D_{client} was lowest, an offer applied to all potential trading partners in more than half of the cases.

The dissemination of offers among competitors provided a completely different pattern. On average, only 12% of the competitors were informed about an offer, and the median value is zero. In 85% of all cases not one competitor was informed. It is evident that subjects were very reluctant to share information about their offers with their competitors.

A similar picture arises if one looks at the dissemination of asks and bids separately. Fig. 2 shows the relative frequencies of the different levels of D_{client} and D_{comp} . In most cases an ask applied to all buyers and no other seller was informed

⁶Note that this measure underestimates the actual the dissemination of an offer among potential clients, since sometimes traders knew that some potential trading partners had already left the market.

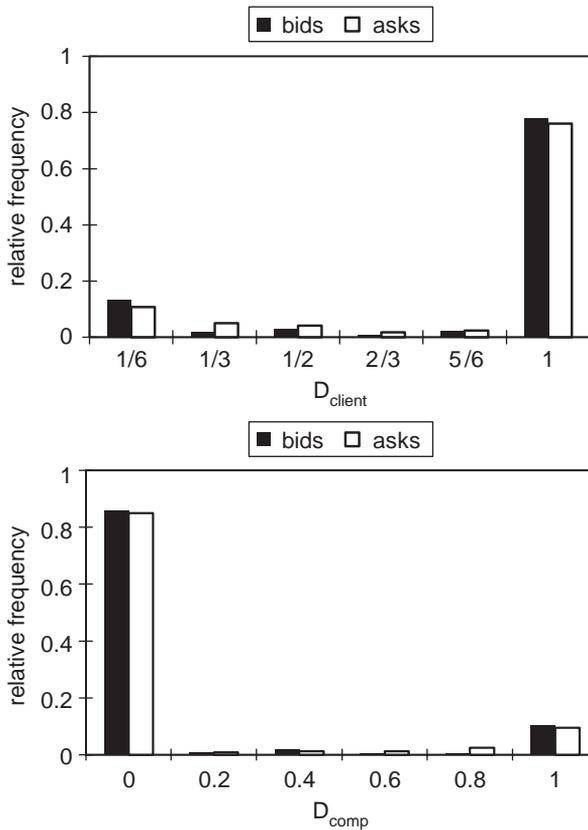


Fig. 2. Relative frequencies of dissemination of bids and asks among potential trading partners (D_{client}) and competitors (D_{comp}) in the DBAM treatment.

about the ask, whereas a bid typically applied to all sellers while no other buyer was informed about it.

This pattern was stable over the time, as can be seen in Fig. 3 that displays the development of average D_{client} and D_{comp} over the 18 rounds. From the very beginning of a session, subjects typically made offers to all potential trading partners and did not inform competitors, and this behavior did not change much over time. We also checked for the development of D_{client} and D_{comp} within the rounds. We found no distinctive differences between the beginning and the end of a round.

Result 2. Prices. *Prices in the DBAM-treatment were as close to the equilibrium price as those in the DAM-treatment and closer to the equilibrium price than those in the DBM-treatment.*

Remember that in our design the equilibrium price was set-valued (see Fig. 1). Therefore, we looked at the average absolute difference between actual prices and

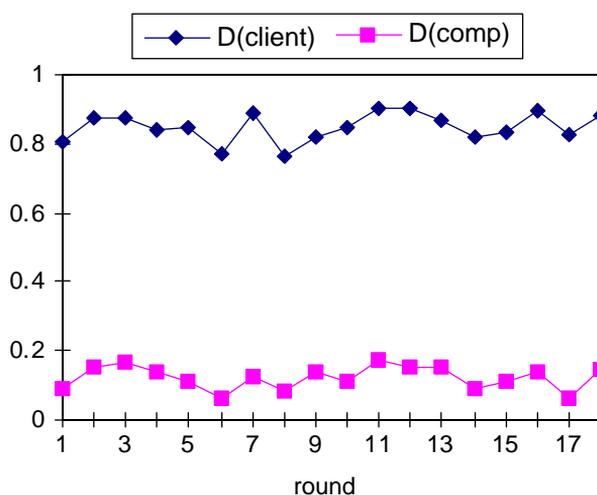


Fig. 3. Development of average D_{client} and D_{comp} in the DBAM-treatment.

Table 3

Average absolute difference between actual prices and equilibrium price range

Session #	DAM	DBAM	DBM
1	0.57	0.54	2.04
2	1.72	0.88	1.26
3	1.49	0.67	1.76
4	0.61	1.08	1.31
All sessions	1.1	0.79	1.59

the equilibrium price range [30,35]. The endogenous market with no frictions has an average distance to the equilibrium price that is not significantly different from the double auction market ($p = 0.68$).⁷ On the other hand, the distance to the equilibrium price was larger in the DBM-treatment (1.59) than in the DBAM-treatment.

As can be seen in Table 3, even prices in session DBAM4 (the DBAM-session with the largest distance between actual and equilibrium prices) were closer to equilibrium than prices in DBM2 (the DBM-session with the lowest distance to equilibrium). Hence, the difference between DBAM and DBM was highly significant ($p = 0.028$).

One purpose of market institutions is to provide a matching and information structure that allows the traders to reap the gains of trade. Did the endogenously emerging market institution serve that purpose? We investigated the gains of trade not reaped by the traders, i.e., the extent to which the sum of consumer and producer surplus fell short of the maximum level. One measure of inefficiency is the foregone surplus as percentage of the maximum possible surplus. Another measure is the

⁷Throughout we employed two-tailed Mann–Whitney tests with session averages as observations.

Table 4
Average levels of (a) foregone surplus and (b) the number of inefficient rounds

Session #	DAM	DBAM	DBM
(a)1	1.48	1.77	6.31
2	0.25	0	7.83
3	1.26	1.26	9.34
4	1.77	0.76	5.30
All sessions	1.19	0.95	7.20
(b)1	3	3	13
2	1	0	11
3	3	4	11
4	4	2	6
All sessions	11	9	41

number of inefficient rounds, i.e., the number of rounds in which the foregone surplus was larger than zero.

Result 3. Efficiency.

- (a) *The foregone surplus in the DBAM-treatment was as low as in the DAM-treatment.*
- (b) *The foregone surplus in the DBAM-treatment was significantly lower than in the DBM-treatment.*

For each of the three treatments, Table 4 presents the average levels of foregone surplus and the number of inefficient rounds.

The average inefficiency levels in the DBAM and the DAM sessions were about 1% and in the DBM sessions it was about 7%.⁸ (see Table 4). The sealed offer double auction market that endogenously emerged in the DBAM-treatment was about as efficient as the DAM. The efficiency levels of the individual sessions of the two treatments overlapped and the difference between the treatments was not significant ($p = 0.68$). This shows that an endogenous market institution can be as efficient as the DAM. The least efficient institution was the DBM. Even the most efficient DBM-session (DBM4) was less efficient than the least efficient DBAM-session (DBAM1). Again, the difference between the DBAM and the DBM-treatment was highly significant ($p = 0.028$).

Qualitatively the same results appeared when we looked at the number of inefficient rounds (Table 4). In the DBAM (DAM) treatment only 9 (11) of 72 periods were inefficient. In the DBM in more than half of the rounds not all gains of trade were reaped (41 inefficient rounds out of 72). Again the difference between DBAM and DBM was significant ($p = 0.028$), whereas the difference between DAM and DBAM was not ($p > 0.5$).

⁸The impact of a trade of an extra-marginal trader as well as that of a non-trade of an intra-marginal trader on efficiency depends of course on the supply and demand conditions. Hence, not the absolute values of foregone surplus but the differences in efficiency between the treatments are important.

For a further examination of the efficiency properties one can distinguish between two different types of inefficiencies. A round is inefficient whenever an extramarginal buyer or seller trades. Such inefficient trades can only occur at out-of-equilibrium prices. But even in the absence of inefficient trades, a round may be inefficient if not all intramarginal traders strike a deal. Such an inefficient nontrade can occur if potential trading partners do not find each other or if they cannot agree on the price. It turned out that the inefficiencies during the early rounds of a session were due to inefficient trades as well as inefficient non-trades. If a later round was inefficient, however, this was almost always due to inefficient non-trades. This pattern emerged in all three treatments. Both types of inefficiencies, however, were much more frequent in the DBM than in the DBAM- and the DAM-treatment.

4. Transaction costs

In the previous section we studied the market institution that emerges when there are no exogenous frictions. In this section we study the effect of exogenous transaction costs. In particular, we examine how the market institution adapts to the introduction of buyer communication costs. For an overview of the treatments see Table 1.

4.1. Treatment 4: Directed bid–ask market with transaction costs (BDAM–TC)

The market technology was as in the directed bid–ask market. Whenever traders made an offer they decided to which potential trading partners and to which competitors the offer should be directed. Sellers could still submit asks at no cost. However, submitting bids to the market was now costly for the buyers. For each submitted bid, a buyer incurred a cost of 0.25 points for every trader (buyer or seller) he chose to inform of the offer.

In this treatment we expected buyers to be less active by making both less bids and possibly also broadcasting those bids less widely. Therefore, the emerging institution should be rather one-sided with sellers doing most of the active offering, while buyers mainly rely on waiting and accepting offers (which is costless to them). Given that traders chose not to inform competitors in the DBAM-treatment, we expected sellers to exhibit the same behavior in the BDAM-TC treatment.⁹ Together these choices imply that there would be very little information to the active market side, the sellers. This may lead to inefficiencies. There exists, however, an institution that is as efficient as the double auction market, and entails no transaction costs: the offer auction (Walker and Williams, 1988). In an offer auction only sellers submit offers (asks). These asks apply to all buyers and every trader is informed of all asks and all transaction prices. To achieve this institution, sellers would have to inform each other about the offers they make, that is, to communicate their prices.

⁹Buyers now incur a cost for informing competitors. We do not expect that this will lead them to inform their competitors more often.

4.2. Results for the DBAM-TC

Result 4. Market institution. *The institution emerging in the DBAM-TC treatment is best described as a sealed offer auction. Most of the offers (asks) were made by sellers, and these offers were sent to every buyer but concealed from competitors.*

The numbers of asks and bids in the DBAM-TC treatment provide first evidence for Result 4. Whereas the introduction of the one-sided transaction costs left the average number of asks per round unchanged, it decreased the average number of bids from 15.7 to 8.6 (see Table 5).

The effect of transaction costs on the dissemination of bids among sellers was even more dramatic than the effect on the number of bids. While in the DBAM treatment the average D_{client} of bids was 0.858, it dropped to a level of 0.384 in the DBAM-TC treatment (see Table 6).

The dissemination of asks did not change with the introduction of the transaction costs. Average D_{client} in the DBAM-TC treatment was 0.853, whereas in DBAM it was 0.865—typically asks applied to essentially all buyers (see Table 6). When we combined the effect of the decrease in the number of bids with the effect of the lower dissemination of bids, we found that the average number of bids a seller received per round dropped from 13.4 in the DBAM to 3.3 in the DBAM-TC treatment. On the other hand, the average number of asks a buyer receives per round did not change (14.3 in the DBAM, 15.1 in the DBAM-TC treatment), and was about 5 times as large as the number of bids a seller received in the DBAM-TC treatment. Hence, we can conclude that the introduction of one-sided transaction costs induced the emergence of an institution in which offers (asks) were made by the sellers, but not by the buyers. Furthermore, the behavior of traders was quite stable across periods. The average number of bids in the DBAM-TC treatment declined slightly during the course of the experiment, indicating that with experience the buyers were somewhat less willing to make bids.

Table 5
Average number of offers per round

Offer	Session	DBAM	DBAM-TC
Bids	1	14.4	11.2
	2	15.1	5.6
	3	13.0	9.7
	4	20.3	8.1
	Average	15.7	8.6
Asks	1	20.6	21.6
	2	16.4	15.2
	3	11.8	20.4
	4	17.3	13.7
	Average	16.5	17.7

Table 6
Average values of D_{client} and D_{comp}

Offer	Session	D_{client}	D_{client}	D_{comp}	D_{comp}
		DBAM	DBAM-TC	DBAM	DBAM-TC
Bids	1	0.782	0.337	0.099	0.009
	2	0.875	0.411	0.080	0.039
	3	0.955	0.423	0.091	0.016
	4	0.820	0.366	0.203	0.004
	Average	0.858	0.384	0.118	0.017
Asks	1	0.627	0.838	0.023	0.117
	2	0.965	0.810	0.101	0.042
	3	0.923	0.880	0.090	0.074
	4	0.945	0.884	0.288	0.122
	Average	0.865	0.853	0.125	0.089

Like in the DBAM treatment, the dissemination of offers among competitors (D_{comp}) was very low in the DBAM-TC treatment (see Table 6). Hence, in both treatments offers were kept secret from the competitors.

Result 4 supports the idea that transaction costs are capable of shaping the behavior of market participants and the resulting market institution. The introduction of one-sided transaction costs led to a sharp decrease of costly bids, and the few remaining bids were less widely disseminated.

What is the information structure of the emerging market institution? In a ‘pure’ sealed offer auction, only asks are made and competing sellers are never informed. As a consequence, every seller knows only his own asks and his own realized trades. A seller would have even less information than in the DBM, where he sometimes observes bids of buyers. Of course, in the DBAM-TC sessions the emerging sealed offer auction was ‘impure’, i.e. some bids were made and some asks were transmitted to other sellers. In the DBAM-TC treatment an average seller was informed about 19.8% of the offers made by others, and about 10.8% of the acceptances of offers made by others. In comparison, in the DBAM treatment an average seller was informed about 51.5% of the offers made by others, and about 50.2% of the acceptances of offers made by others. By design, in the DBM treatment these numbers are 9% and 0%, and in the DAM both numbers are 100%. Furthermore, if we look at the development over the periods, we observe no trend towards more information provision for sellers. The percentage of offers a seller is informed of even decreases slightly as buyers learn to avoid transaction costs by entering fewer bids.

Next we investigate efficiency properties.¹⁰ We have already seen that buyers do incur some transaction costs. We now investigate the efficiency of the transactions (gains from trade), not taking into account the wasteful transaction costs.

¹⁰There were no significant differences between the DBAM-TC treatment and the DBAM treatment with respect to the difference between actual prices and the equilibrium price range.

Table 7

Average levels of (a) foregone surplus and (b) the number of inefficient rounds

Session #	DAM	DBAM	DBAM-TC	DBM
(a)1	1.48	1.77	3.54	6.31
2	0.25	0	2.02	7.83
3	1.26	1.26	3.79	9.34
4	1.77	0.76	2.02	5.30
All sessions	1.19	0.95	2.84	7.20
(b)1	3	3	5	13
2	1	0	5	11
3	3	4	7	11
4	4	2	5	6
All sessions	11	9	22	41

Result 5: Efficiency with transaction costs. *The foregone surplus was significantly higher in the DBAM-TC treatment than in the DBAM treatment and the DAM. However, the foregone surplus was significantly lower than in the DBM.*

The average inefficiency level was about 3% in the DBAM-TC treatment as compared to 1% in the DAM and DBAM treatments and 7% in the DBM treatment (see Table 7a). Even the most efficient DBAM-TC sessions (DBAM-TC2 and DBAM-TC4) were less efficient than the least efficient DBAM (DBAM1) and DAM (DAM4) sessions. The difference between DBAM-TC and these two other treatments was highly significant ($p = 0.028$). However, the least efficient DBAM-TC session (DBAM-TC3) was more efficient than the most efficient DBM session (DBM4), and efficiency differences were significant at ($p = 0.028$).

A similar result holds when we consider the number of inefficient rounds (Table 7b). These inefficiencies can be caused by inefficient trades (at out of equilibrium prices) or by inefficient non-trades. Both types occur in early rounds, whereas in later rounds inefficiencies are mostly caused by inefficient non-trades.

5. The effects of price communication

So far the emerging markets in our experiments share the feature that traders conceal offers from their competitors. When, as in the sealed offer double auction market, both sides of the market were active the institution nevertheless led to high levels of efficiency. However, when only one side of the market was active, as in the sealed offer auction, the resulting institution was inefficient. Furthermore, the efficiency of offer auctions suggests that this preference of traders to conceal offers is the culprit for the observed inefficiency.

To test for the robustness of this preference for concealing offers, and to examine the effects of price communication between competitors on market outcomes, our last treatment provides traders with explicit incentives to communicate prices to

competitors. Specifically traders have to incur (small) costs if they want to conceal offers from their rivals. This will show whether traders (sellers) have a weak or indeed a strong preference for not informing their competitors. Furthermore, if sellers chose to inform each other, we can study the effect of price communication on prices and efficiency.

5.1. Treatment 5: Directed bid–ask market with subsidized price communication (DBAM-PC)

In treatment 5 (DBAM-PC), just as in treatment 4 (DBAM-TC), buyers bear transaction costs of 0.25 points for each trader (buyer or seller) they choose to inform of an offer. Furthermore, a seller incurs a fixed cost of 0.02 points for each competitor she does not inform of an offer. The options to conceal or reveal price quotes are identical to those in treatment 4, only the direct costs of doing so are different. For example, a price quote can be concealed from all five competitors at zero costs in treatment 4, whereas doing so entails a cost of 0.1 points (5×0.02) in treatment 5.

The directed bid–ask market with transaction costs (DBAM-TC) now serves as a baseline treatment to measure the preferences of sellers to conceal offers, and to study the effects of price communication (for an overview of all treatments see [Table 1](#)).

5.2. Expected results for the DBAM-PC

The results of the endogenous market with subsidized price communication will show whether traders (sellers) have a weak or indeed a strong preference for not informing their competitors. What are the possible gains of sellers when they choose to inform each other?

First, sellers would avoid the costs for concealing offers from competitors. Second, we have seen that the DBAM-TC treatment led to inefficiencies, presumably because the active market side (the sellers) did not receive enough information. If sellers would choose to always inform each other, then the resulting institution would resemble an offer auction. In an offer auction only sellers make asks. These asks apply to all buyers and every trader is informed about all asks and about every acceptance of an ask. From a previous study ([Walker and Williams, 1988](#)) we know that an offer auction exhibits the same properties as the DAM. Therefore, price communication among competitors might indeed lead to higher efficiencies and sellers may be able to claim part of this surplus.

Third, price communication may raise prices. Remember that in our market ([Fig. 1](#)) the competitive quantity is 4 and the equilibrium price range is [30,35]. At a market price of 35, the joint profit of sellers is 70 (and buyers' profits are 40), which is a considerable improvement over the joint profits of 50, which they attain at a price of 30 (where buyers' profits will be 60). Therefore, more information among competitors might lead to a higher average transaction price than the one (31.5) achieved in the treatment without a price communication subsidy (DBAM-TC).

5.3. Results for the DBAM-PC

Result 6: Communication. *Most sellers strictly prefer to conceal their offers from rivals.*

In a majority of the cases (59.6% averaged across sessions) sellers do not inform any competitor even if this secrecy is costly. In the remaining cases at least one rival is informed about the offer. In those cases, typically *all* rivals are informed (averaged across sessions we have $D_{\text{comp}} = 1$ for 38.4% of the asks). The average dissemination of offers is higher in the endogenous market with subsidized price communication (DBAM-PC) than in that with transaction costs (DBAM-TC), but at $D_{\text{comp}} = 0.377$ it is still far below unity. Many sellers are willing to pay for secrecy. There is no tendency for D_{comp} to change much over time.

As can be seen in Table 8, there is quite some heterogeneity across sessions. Whereas in session 2 almost all sellers opted to conceal prices, in DBAM-PC4 sellers quite often chose to inform each other.

Result 7: The effects of prices communication. *Price communication between sellers has a small positive effect on transaction prices.*

Table 9 presents average transaction prices for the different sessions and treatments. It appears that average prices are higher in the DBAM-PC treatment, (32.6) than in DBAM-TC treatment (31.5). The difference is small and insignificant ($p=0.11$) with a two-tailed Mann–Whitney test and the 8 session averages as observations. However, if we restrict attention to the first 12 rounds, average prices in treatment DBAM-TC are persistently below those in treatment DBAM-PC, and the difference is significant ($p=0.028$). Hence, as long sellers as cannot rely too much on the experiences from past rounds, increased price communication between sellers has an upward effect on transaction prices.

Since there is a lot of heterogeneity on the amount of price communication across sessions, we analyze the relation between sellers' communication and transaction prices in the DBAM-PC treatment at a more disaggregate level. For each round we

Table 8
Dissemination of ask prices among sellers

Session #	Average value of D_{comp}		% offers with $D_{\text{comp}} = 0$	
	DBAM-TC	DBAM-PC	DBAM-TC	DBAM-PC
1	0.117	0.394	85.4	58.0
2	0.042	0.016	95.2	97.8
3	0.074	0.466	89.5	47.4
4	0.122	0.632	85.8	35.3
Average	0.089	0.377	89.0	59.6

Note: Averages and percentages are taken first over the offers within a round and then over the rounds.

Table 9

Mean and standard deviation of transaction prices in the Directed Bid–Ask Market with Transaction Costs (DBAM-TC) and with Price Communication (DBAM-PC).

Session #	Mean price		Standard deviation	
	DBAM-TC	DBAM-PC	DBAM-TC	DBAM-PC
1	32.6	33.4	2.30	3.43
2	31.8	31.4	2.87	3.58
3	31.1	32.8	4.99	3.70
4	30.3	33.0	4.26	2.56
Average	31.5	32.6	3.60	3.32

Note: Means and standard deviations of prices are taken first over the offers within a round and then averaged over the rounds.

calculate the average level of D_{comp} across all asks (that is, for accepted as well as unaccepted asks). This gives us a measure for the general level of sellers' price communication in a round. We compare this to the corresponding average value of the transaction prices in the round. Across all rounds and sessions, the Pearson correlation between these two values is significantly positive: $r_p = 0.44$ ($n = 72$, $p < 0.01$).

The impact of price communication may arise from the fact that offers, which were concealed from rivals ($D_{\text{comp}} = 0$), were significantly lower ($p = 0.066$) than offers that were revealed to at least one other seller ($D_{\text{comp}} > 0$).¹¹

In sum, we find that increased price communication between sellers increases transaction prices. The effect is statistically significant only in the early rounds and moderate in absolute terms. In particular, average prices typically stay within the range defined by the competitive equilibrium.

We have seen that sellers display a preference to conceal price quotes (Result 6). Furthermore, Result 7 indicates that sellers benefit collectively from revealing price quotes. The next result shows that revealing price quotes has a public goods characteristic.

Result 8. *It is individually rational to conceal price offers from competitors.*

For each session in the DBAM-PC treatment, Table 10 reports the average price of accepted asks. A distinction is made between accepted asks concealed from rivals ($D_{\text{comp}} = 0$) and accepted asks revealed to at least one rival ($D_{\text{comp}} > 0$). We find the accepted asks with $D_{\text{comp}} = 0$ have an average price of 33.7 and those with $D_{\text{comp}} > 0$ have an average price of 33.3.¹² Although the difference is small, it is significant at a

¹¹The average price of offers with $D_{\text{comp}} > 0$ was 42.2, compared to 38.3 for $D_{\text{comp}} = 0$.

¹²Of all accepted asks in the DBAM-TC treatment, those with $D_{\text{comp}} > 0$ have an average price of 30.9 and those with $D_{\text{comp}} = 0$ an average price of 32.1. There are however only 9 accepted asks with $D_{\text{comp}} > 0$ versus 180 with $D_{\text{comp}} = 0$. Therefore, the table focuses on the sessions of the DBAM-PC treatment.

Table 10
Average price for concealed and revealed accepted offers

Session #	$D_{\text{comp}} = 0$	$D_{\text{comp}} > 0$
1	34.8	33.0
2	32.4	—
3	33.6	33.5
4	34.2	33.3
Average	33.7	33.3

Note: $p = 0.08$ Wilcoxon matched-pairs signed-ranks test ($n = 3$).

10% level ($p = 0.08$, $n = 3$). Furthermore, a price difference of 0.4 is sufficient to warrant the fine of concealment ($0.1 = 5 \text{ rivals} \times 0.02$) if no more than four price quotes are made before acceptance (the average number of quotes per seller per round is 1.6).

These data suggest that concealing price quotes from rivals is indeed the individually rational thing to do. Even though sellers benefit collectively from price sharing, an individual seller has little interest to reveal its prices to competitors. By revealing price quotes to competitors, a seller is more prone to being undercut by a rival seller. Concealed offers can carry a higher price than revealed offers.

We now explore the effects of price communication on efficiency.¹³

Result 9. *Low levels of price communication did not significantly improve market efficiency, compared to the treatment without subsidized price communication.*

Table 11 indicates that the average level of efficiency is higher in treatment DBAM-PC than in treatment DBAM-TC, though not as high as in treatment DBAM. None of the differences are statistically significant, however. In DBAM-PC sessions 1–3, the sessions with low levels of price communication, the efficiency was as low as in treatment DBAM-TC. Only session DBAM-PC4, in which sellers informed each other of about 65% of the offers, showed a higher level of efficiency, comparable to the efficiency in treatment DBAM.

The results of the market with subsidized price communication suggest that many traders do not just have a weak, but a strong preference to conceal information from their competitors. At the same time, if traders would manage to disclose information to competitors they could benefit collectively. Price communication among competitors is thus characterized by public good features. Even though sellers would collectively benefit, each individual seller has an incentive to keep her own ask secret.

¹³We note furthermore that price communication had no effect on price dispersion or on the speed of convergence. This should not be surprising since prices in the DBAM-TC treatment were already quite close to equilibrium prices, and neither farther from equilibrium than in the DBAM treatment, nor closer to equilibrium than in the DBM treatment.

Table 11
Foregone surplus (in % of total surplus)

Session #	DBAM	DBAM-PC	DBAM-TC
1	1.77	2.8	3.54
2	0	2	2.02
3	1.26	3	3.79
4	0.76	0.3	2.02
All sessions	0.95	2.02	2.84

6. Conclusion

To investigate the forces that shape market institutions we conducted experiments in which the information and matching process was endogenized. In our baseline design the resulting institution can be best characterized as a ‘sealed offer double auction market’ as traders typically chose to conceal offers from their competitors. This institution provides less information to traders than the double auction market (in which traders learn about all offers and all transaction prices). Nevertheless, the sealed offer double auction market mimicked the double auction market in terms of efficiency and convergence to market clearing prices. In contrast, in a decentralized bargaining market, in which each trader is only allowed to communicate with one potential trading partner at a time, prices converged more slowly and efficiency was lower.

When imposing one-sided transaction costs on the buyers, the sellers made most of the offers. This shows that even small transaction costs can shape the matching and information structure of markets. The preference of sellers to conceal their offers resulted in a sealed offer auction market which led to inefficient trading outcomes.

When sellers had to bear small costs to hide information from competitors, many sellers were willing to incur these costs. Providing information to competitors has features of a public good. Collectively the sellers would profit, by maintaining higher prices, though individually each seller has an incentive to keep her prices secret.

In the present paper we have focussed on a basic environment with a homogenous good, one unit per trader, and stationary aggregate demand and supply. Future work can examine the market institution that emerges in more complex environments. For example, it would be interesting to see whether the presence of market power will affect the incentives of sellers to communicate price offers. Also, it would be interesting to examine whether the positive effect of seller communication on prices that we found only in the earlier periods of our stationary environment, will persist for longer if there are random shocks to supply and demand. More generally, by varying the characteristics of the market environment we can get a better understanding of the forces that shape the market institution and of the corresponding properties of this institution in terms of prices, convergence and efficiency.

Acknowledgements

We are grateful to Joep Sonnemans for programming (using RatImage toolbox, Abbink and Sadrieh, 1995), to an anonymous referee, and to the participants of seminars at the universities of Amsterdam, Berlin, Bonn, Copenhagen, Harvard, Helsinki, Karlsruhe, Leuven, Queen Mary, Rotterdam, Stockholm, Tilburg, Vienna, Zürich, the Review of Economics Studies tour, a workshop in Koblenz, the ESA conference in Mannheim, and the ESA-ESEM conference in Venice for helpful advice. The research by J. Potters is made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences. Financial support by the EU financed TMR-project on “Savings and Pensions” (Contract No. ERBFMRXCT 960016) is highly appreciated. The usual disclaimer applies.

Appendix

This appendix contains the instructions for the DBAM-TC treatment. For the treatments DAM, DBM, DBAM and DBAM-PC we adapted the instructions accordingly.

Instructions

Today we are going to set up a market in which some of you will be buyers and others will be sellers. Those of you who have drawn a “B” are buyers, those of you with an “S” are sellers. During the experiment you will have to make trading decisions and these decisions will determine your earnings. During the experiment your earnings will be denoted in points. For each point you earn you will get 30 cents for your participation. Your earnings will be paid to you privately and in cash after the experiment.

First, we will go through the instructions together. After that you will get the opportunity to study the instructions in your own pace and to ask questions. Also we will have a practice round before we start the experiment. If you want to make notes you can use the empty sheet that is on your table. Please, do not write on the instructions, and do not touch the keyboard until we are done with the instructions.

The experiment will consist of 18 trading rounds. In each trading round each buyer may buy at most one unit. Each seller may sell at most one unit. Trades are made in accordance with certain rules that will be explained below. If a buyer buys a unit in a round, her gross-earnings are equal to the *resale-value* of the unit to her minus the price she pays for the unit. If a buyer does not buy a unit, her gross-earnings for that round are zero. If a seller sells a unit in a round, his gross-earnings are equal to the price he receives for the unit minus the *cost-value* of the unit to the seller. If a seller does not sell a unit, his gross-earnings for that round are zero.

Resale-values as well as cost-values are strictly private information, no other participant than the concerning buyer or seller, respectively, will learn about it.

Hence, every participant knows only his own gross-earnings from a trade. Your net-earnings in a round are equal to your gross-earnings minus the cost you made by sending offers to other traders. How these latter costs are determined will be explained below. Finally, your total earnings for the experiment are equal to the sum of your net-earnings in each of the 18 rounds.

Computer screen

The trading of units will take place by means of the computer. All relevant information will be available on your computer screen. You can now see what the screen will look like during the experiment.

In the top left of the screen you can see how many points you have earned up to that moment. Also the number of the present round is indicated. Below that you see how much time is left for trading in that round. For each of the 18 rounds the total time for trading is 3 minutes.

In the middle of the screen you will see a list of the relevant ask and bid-prices. An ask-price indicates the price at which a seller is prepared to sell, and a bid-price indicates the price at which a buyer is prepared to buy. Ask- and bid-prices will be explained in more detail below.

On the right-hand side of the screen your value is indicated. If you are a buyer, the *resale-value* of a unit in that round is indicated. If you are a seller, the *cost-value* of a unit in that round is indicated. Resale-values may be different for different buyers, and cost-values may be different for different sellers. Also your value may change from round to round. But every buyer gets the same resale-values the same number of rounds, and every seller gets the same cost-values the same number of rounds.¹⁴ Also you can see where your gross-earnings will be indicated, and that your net-earnings for the round will be equal to your gross-earnings minus the cost you have made in sending offers (as will be explained below).

In the middle left of the screen you see whether you are a buyer or a seller. Here also your ID-letter is indicated. The IDs for the buyers are A, B, C, D, E and F. The IDs for the sellers are U, V, W, X, Y and Z. Your ID will randomly change from round to round. Furthermore, IDs are not related to the cost- or resale-values. Hence, the IDs are for registration purposes only and do not convey information about the identity of participants or about their values.

In the bottom right of the screen you see a box called last action. This box mentions the last action that is relevant to you.

Finally, at the bottom of the screen you see a long flat row. In this row you are informed about all accepted asks and bids that were in your column of ask- and bid-prices.

Now we will first go through the specific instructions for buyers, and then through the specific instructions for sellers. On your handout you will only find the specific instructions that concern your role in the market.

¹⁴Although this is not in the written instructions, at this point we told the subjects in all treatments that the set from which the cost and resale-values were drawn was constant across rounds.

Specific instructions for buyers

After the determination of the values and the IDs, the market opens for trading. If you want to buy a unit, and you have not yet bought a unit in that round, then you can do two things:

(1) You can press **B** to make a *bid* and to enter a price at which you are prepared to buy a unit. After you press **B** you are requested to enter a bid-price. This bid-price must be above or equal to zero and below or equal to your resale-value. After you have entered your bid-price, you must decide to which sellers you want to send your bid and which of the other buyers you want to inform about your bid. *You may send your bid to any number of sellers, but at least to one seller. You may inform any of the other buyers about your bid, but you are not obliged to do so.* Therefore, after you entered your bid-price, you are also requested to enter the ID(s) of at least one seller to whom you want to send this bid and of any of the other buyers you want to inform. Hence, you may enter as many of the seller ID-letters (U, V, W, X, Y, Z) and buyer ID-letters (A, B, C, D, E, F) as you want, but you should at least enter one seller ID-letter. Recall that the IDs change from round to round. After you have entered the ID(s), your bid-price will appear in the lower middle of your screen in the column “bids” and it is marked with an asterisk (*). Now your bid-price is transmitted to the sellers and buyers you have entered, and will appear in their columns of bid-prices together with your ID.

By sending your bid to a seller or another buyer you incur a cost. To be precise, for each of the other traders (seller or buyer) that you decide to inform about your bid, you incur a fixed cost of 0.25 points. The total cost you incur by sending bids to other traders will be subtracted from your gross-earnings to determine your net-earnings for the round.

A seller who receives your bid may accept it or not. The buyers you selected can of course not accept your bid—they are only informed about it. As soon as one of the sellers accepts your bid, you will get a message in the lower-right corner of your screen under “last action”, and the corresponding earnings will be calculated and indicated on your screen under “value”. Also all other sellers and buyers you have chosen to send your bid to will learn that your bid is accepted. Your bid-price will vanish from their column “bids” and will appear in the row at the bottom of their screens.

If your bid is not accepted by a seller, you will not get a message. Notice that it is possible that you send a bid to a seller who has already sold a unit to another buyer. Hence, receiving no message may mean that the sellers you selected to send your bid to have not yet decided about your bid, that they reject it, or that they already sold a unit.

As long as you have not bought a unit in the trading round you may always revise your bid by just pressing “**B**”, entering a (possibly new) price, and entering (possibly new) sellers’ and buyers’ IDs. If you enter a new bid, your old bid will disappear. By making a new bid you again incur a fixed cost for each of the other traders that you decide to inform about your new bid. However, you are not forced to make any (new) bid. Like buyers can make bids, sellers can make *asks*. This brings us to the second thing you can do to buy a unit.

(2) You can press on K. In that case you can buy one unit at one of the ask-prices that is indicated in the column “asks” in the upper-middle of your screen. Of course, you cannot use this option if this column is empty. If an ask-price is indicated in this column, then one of the sellers (indicated by his ID-letter) is prepared to sell a unit to you at the indicated price. If you decide to accept this ask-price you simply press K. Since there may be more than one ask-price in the column, you also need to enter the seller-ID of the ask you wish to accept. Then the trade is conducted, your earnings are registered, and all buyers and sellers who were informed about this ask can see at the bottom of their screen that a trade has occurred at this price.

Specific instructions for sellers

After the determination of the values and the IDs, the market opens for trading. If you want to sell a unit and you have not yet sold a unit in that round, you can do two things:

(1) You can press P to make an *ask* and to enter a price at which you are prepared to sell a unit. After you press P you are requested to enter an ask-price. This ask-price must be below or equal to 100 and above or equal to your cost-value. After you have entered your ask-price, you must decide to which buyers you want to send the ask and which of the other sellers you want to inform about your ask. *You may send your ask to any number of buyers, but at least to one buyer. You may inform any of the other sellers about your offer, but you are not obliged to do so.* Therefore, after you entered your ask-price, you are also requested to enter the ID(s) of at least one buyer to whom you want to send this ask, and of any of the other sellers that you want to inform. Hence, you may enter as many of the buyer ID-letters (A, B, C, D, E, F) and seller ID-letters (U, V, W, X, Y, Z) as you want, but you should at least enter one buyer ID-letter. Recall that the IDs change from round to round. After you have entered the ID(s), your ask-price will appear in the upper middle of your screen in the column “asks” and it is marked with an asterisk (*). Now your ask-price is transmitted to the buyers and sellers you have entered, and will appear in their columns of ask-prices together with your ID.

For each of the other traders (buyer or seller) that you decide to inform about your ask, you incur a fixed cost of zero—by sending your ask to a buyer or another seller you does not incur any cost. Hence if you make a trade in a round, your net-earnings for that round are equal to your gross-earnings.

A buyer who receives your ask may accept it or not. The sellers you selected can of course not accept your ask—they are only informed about it. As soon as one of the buyers accepts your ask, you will get a message in the lower-right corner of your screen under “last action”, and the corresponding earnings will be calculated and indicated on your screen under “value”. Also all other buyers and sellers you have selected to send your ask to will learn that your ask is accepted. Your ask-price will vanish from their column “asks” and appear in a row at the bottom of their screens.

If your ask is not accepted by a buyer, you will not get a message. Notice that it is possible that you send an ask to a buyer who has already bought a unit from another seller. Hence, receiving no message may mean that the buyers you selected to send

your ask to have not yet decided about your ask, that they reject it, or that they have already bought a unit.

As long as you have not sold a unit in the trading round you may always revise you ask by just pressing “P”, entering a (possibly new) price, and entering (possibly new) buyers’ and sellers’ IDs. If you enter a new ask, your old ask will disappear. However, you are not forced to enter any (new) ask. Like sellers can send asks, buyers can make *bids*. This brings us to the second thing you can do to sell a unit.

(2) You can press on V. In that case you can sell a unit at one of the bid-prices that is indicated in the column “bids” in the lower middle of your screen. Of course, you cannot use this option if this column is empty. If a bid-price is indicated in this column, then one of the buyers (indicated by her ID-letter) is prepared to buy a unit from you at the indicated price. If you decide to accept this bid-price you simply press V. Since there may be more than one bid-price in the column, you also need to enter the buyer-ID of the bid you wish to accept. Then the trade is conducted, your earnings are registered, and all sellers and buyers who were informed about this bid-price can see at the bottom of the screen that a trade has occurred at this price.

Summary

The experiment consists of 18 trading rounds, and each round lasts 3 minutes. You are either a buyer or a seller. In a round each buyer may try to buy one unit and each seller may try to sell one unit. For a buyer, gross-earnings will be equal to the resale-value of the unit minus the price paid. For a seller, gross-earnings will be equal to the price received minus the cost-value of the unit. Values are different for different traders, and they change from round to round. Buyers can try to buy by making bids or by accepting asks. Sellers can try to sell by making asks or accepting bids. Buyers’ net-earnings in a round are equal to their gross-earnings minus the cost they made in sending bids to other traders. Each time a buyer makes a bid she incurs a cost of 0.25 points for *every* seller or buyer she decides to inform about this bid. Sellers costs from sending asks to other traders are zero. Therefore, their net-earnings are equal their gross earning.

During the experiment all earnings are denoted in points. After the experiment, your earnings in cash will be determined at a rate of 1 point = 30 cents. You will receive your earnings privately, immediately after the experiment. Your earnings are your own business, you do not have to discuss them with anyone.

Final remarks

During the experiment, it is not allowed to talk or communicate with other participants in any way (other than through the trading). If you have a question, please raise your hand and the experimenter will come to your table. If anything strange appears on your screen, or if you think the computer is not doing what you think it should, please notify the experimenter so he can try to fix the problem.

References

- Abbink, K., Sadrieh, A., 1995. RatImage—Research Assistance Toolbox for Computer-Aided Human Behavior Experiments. Discussion Paper No. B-325 University of Bonn.
- Campbell, J., LaMaster, S., Smith, V., Van, B., 1991. Off-floor trading, disintegration, and the bid-ask spread in experimental markets. *Journal of Business* 64, 495–522.
- Chamberlin, E.H., 1948. An experimental imperfect market. *Journal of Political Economy* 56, 95–108.
- Easley, D., Ledyard, J.O., 1993. Theories of price formation and exchange in double oral auctions. In: Friedman, D., Rust, J. (Eds.), *The Double Auction Market. Institutions, Theories, and Evidence*. Addison-Wesley, Reading, MA.
- Friedman, D., 1991. A simple testable model of double auction markets. *Journal of Economic Behavior and Organisation* 15, 47–70.
- Gjerstad, S., Dickhaut, J., 1995. Price formation in double auctions. *Games and Economic Behavior* 22, 1–29.
- Gode, D.K., Sunder, S., 1997. What makes markets allocatively efficient. *Quarterly Journal of Economics* 102, 603–630.
- Holt, C.A., 1995. Industrial Organization. In: Kagel, J., Roth, A. (Eds.), *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ.
- Plott, C.R., 1982. Industrial organization theory and experimental economics. *Journal of Economic Literature* 20, 1485–1527.
- Smith, V.L., 1962. An experimental study of competitive market behavior. *Journal of Political Economy* 70, 111–137.
- Stigler, G.J., 1964. A theory of oligopoly. *Journal of Political Economy* 72, 44–61.
- Walker, J., Williams, A., 1998. Market behavior in bid, offer, and double auctions. *Journal of Economic Behavior and Organization* 9, 301–314.