

WORKING

PAPERS

Georg Kirchsteiger
Muriel Niederle
Jan Potters

Do Market Institutions Adapt Efficiently to Transaction Costs?

August 2000

Working Paper No: 0006



DEPARTMENT OF ECONOMICS

UNIVERSITY OF VIENNA

All our working papers are available at: <http://mailbox.univie.ac.at/papers.econ>

Do Market Institutions Adapt Efficiently to Transaction Costs?

by

Georg Kirchsteiger

Muriel Niederle

Jan Potters

First version: March 2000

This version : August 2000

Abstract

We study an experimental market with an endogenous institution. In particular, the information and matching structure of the market is determined by the decisions of the individual traders. We examine the effects of the imposition of exogenous transaction (communication) costs. We find that the institution adapts in the predicted direction but it does not do so in a fully efficient manner. Traders incur avoidable transactions costs, and the market outcome does not fully exhaust the possible gains from trade.

Keywords: market institution, transaction costs, efficiency

JEL-classification: C9, D4, L1

G. Kirchsteiger (georg.kirchsteiger@univie.ac.at), University of Vienna, Department of Economics, Hohenstaufengasse 9, A-1010 Vienna, Austria; J. Potters (j.j.m.potters@kub.nl), Tilburg University, Department of Economics, PO Box 90153, 5000 LE Tilburg, the Netherlands; M. Niederle (niederle@kuznets.fas.harvard.edu), Harvard University, Department of Economics, Littauer Center, Cambridge, MA 02138, U.S.A.

We are grateful to Joep Sonnemans for programming (using RatImage toolbox, Abbink and Sadrieh, 1995), and to the participants of the 1999 ESA conference in Grenoble for comments. 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.

1. Introduction

Market exchange does not take place in an institutional void but in a framework that specifies the actions, messages, information and contracting opportunities available to the market participants. These institutional "rules of the game" serve to match potential trading partners, convey and aggregate information about their willingness to buy and/or sell, and allow them to sign binding contracts. The details of these market institutions have not received much attention in the neoclassical (Walrasian) theory of market clearing equilibria. New Institutional Economics and developments in market microstructure theory have put market institutions on the agenda of economic inquiry, however.

In the field an enormous variety of different market institutions can be observed. For example, the process of matching and bargaining at a bazaar is completely different from the rules that govern financial markets; the typical (Dutch) rules at flower auctions are different from the (English) rules at art auctions; real estate is sold both at auctions and by means of direct negotiations. Both empirical and experimental work has shown that the institutional framework within which trading takes place is consequential. It may influence the efficiency of the market outcome, the speed of convergence towards equilibrium, and the distribution of the surplus over the market participants (see e.g., Lusht (1997), and Stoll and Whaley (1990) for empirical evidence, and Holt (1995) and Plott (1982) for experimental evidence).

In view of this diversity of market institutions, an important question concerns the origin and the determinants of these institutional differences. One of the main hypotheses of New Institutional Economics is that market institutions are shaped by transactions costs. It takes resources to search for trading partners, gather information, negotiate the terms of trade, and to monitor and enforce these terms. New Institutional Economics postulates that market institutions provide incentives for trade and that they do so in way that minimizes transaction costs.

In the present paper we put this hypothesis to an experimental test. We examine whether transactions costs are a driving force behind the nature of market institutions. Specifically, we set up a trading environment in which some of the rules of interaction are endogenous, that is, determined by the decisions of the individual traders. We study whether these rules adjust to the introduction of exogenous transaction costs, and, in particular, whether this adjustment is efficient. Efficiency requires, firstly, that avoidable transaction costs are in fact avoided, and, secondly, that the market institution serves its function to inform and match potential trading partners such that potential gains from trade are realized¹.

¹ In order to assess the performance of endogenously emerging market institutions one could of course investigate field data and historical evidence. However, environmental factors like demand and supply structure, product characteristics, transactions costs etc. which are decisive for the assessment of trading outcomes are often not well-documented. These problems, arising from the lack of control and observability can be overcome by an experimental investigation where the experimenter can design and control the environment. Therefore, the properties of trading institutions can be easily evaluated in a

Whether market institutions arise and adjust efficiently is a matter of some debate among institutional economists. Some tend to answer this question affirmatively. Hayek (1967) is one proponent of the proposition that markets, if left to their own, will develop institutional rules which allow them to function efficiently. This "Hayekian" position, which has a clear functionalist flavor to it, can also be discerned in the writings of Klein et al. (1978), and Williamson (1975). Others are much less optimistic, though. For example, North (1990, 1993) points out that institutions will tend to serve the interests of those that have the power to change them, and, furthermore, that the emergence of efficient institutions may be hindered by public goods problems. Also, bounded rationality may prevent efficient institutions from forming (see also Roth 1984). Another possibility is that, once an institution has come into place, the rules may become entrenched, or, at least, that its further evolution may be path-dependent (see also Friedman, 1993). Hence, according to this more "Northian" view, there is no guarantee that evolutionary or competitive forces lead to efficient outcomes.

The purpose of market institutions is to match potential trading partners and to provide them with the necessary information to form efficient contracts. Hence, differences in the information structure and the matching procedure can be expected to have a substantial impact on the performance of different market institutions. This expectation is strongly confirmed by the experimental literature. On the one hand there is the well-known experimental result that the continuous double auction (DA) is a type of trading institution that seems to invariably generate outcomes consistent with the predictions of the competitive model (see Smith 1962, 1964). This institution allows traders to make bids (offers to buy) and asks (offers to sell) and to accept other traders' offers at any moment. Hence, all traders are informed about all outstanding bids and asks as well as about all realized trades. As a consequence, the process of price formation and matching is a public and multilateral process. At the opposite extreme is a market with decentralized matching and private bilateral bargaining. This is the institution originally employed by Chamberlin (1948) in the very first experimental market study (see also the so-called "telephone markets" in Grether and Plott, 1984, and Hong and Plott, 1982). Under this institution traders are restricted to contact other traders only one by one, and the offers, counteroffers and contracts made in these bilateral encounters are not revealed to other traders. This decentralized bargaining market (DBM) was found to generate inefficient outcomes, incongruent with the competitive model (see also Kirchsteiger, Niederle and Potters 1998).

In the present study, we endogenize the information structure and matching procedure of the market. When a trader submits an offer, she decides who will be informed about this offer and to whom it applies. She can choose to inform any number of traders from both the own and the other market side about her offer. If a potential trading partner is informed about an offer, it also applies to her. Hence, if

market experiment.

each trader always chooses to inform every other trader about her offer, the information structure of the market is identical to that of a double auction - every trader is always informed about all outstanding offers and every offer applies to every potential trading partner. In that case the market mechanism is in effect a double auction with a centralized information and matching structure. If, on the other hand, each trader always chooses to inform only one other trader about her offer, price formation and matching are a completely decentralized process as in Chamberlin's experiment.

In order to investigate the impact of transaction costs we employ one-sided transaction costs. Whereas sellers have no transaction costs to bear, a buyer has to bear small, but positive transaction costs for every other trader he informs about an offer she makes. Note that one-sided transaction costs allow in principle for the possibility that trade takes place without traders incurring any costs. This allows for a unambiguous assessment of the efficiency of the market outcome. Furthermore, transaction cost can be avoided by the emergence of a one sided institution, where only sellers make offers, but buyers refrain from making bids. Hence, one sided costs allow for a clear prediction of the institution if its emergence is indeed shaped by the transaction costs.

Our paper is organized as follows. In the next section we describe the experimental design. In Section 3 we present the results, and Section 4 contains a concluding discussion.

2. Experimental Design

Our experimental design consisted of three treatment conditions. In the first treatment, we investigated the endogenous emergence of market institutions in the presence of transaction costs. When making an offer, a trader not only specified the price at which he was prepared to trade, but he also indicated which subset of the other traders he wanted the offer to apply to and he wished to inform about the offer. No transaction costs were imposed on sellers, but a buyer had to bear small, but positive transaction costs for every trader he informed about a bid he made. We called this the transaction costs (TC) treatment. We also employed two control treatments. To control for the effect of the transaction costs, we run experiments where not only the sellers but also the buyers could costlessly make offers and inform other traders about them. This was the no transaction cost (NTC) treatment². Furthermore, we also looked at experiments without transaction costs in which the rules of the double auction were exogenously given (DA treatment). Since the DA has been found to be the most efficient institution in market experiments, this control treatment allows us to assess the efficiency properties of the NTC treatment. Apart from these institutional

²Using such an NTC treatment the endogenous emergence of market institutions with absent transaction costs has already been investigated in detail in a previous study (see Kirchsteiger, Niederle and Potters 1998).

manipulations, our design followed the conventions that have been developed regarding experimental markets (see the Appendix for a set of instructions).

Market Structure

Each experimental market session consisted of a sequence of one practice round and 18 trading rounds. Each trading round lasted three 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).

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

Insert Figure 1 about here

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.

Treatments

(a) Transaction cost (TC) treatment

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 he was prepared to trade. He also had to enter the IDs of those traders he wanted to inform about the offer. 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.

As already mentioned, asks were costless. A buyer making a bid, however, had to pay 0.25 points for every seller to whom the bid applied. Furthermore, the price for informing another buyer about the bid was also 0.25 points. Notice that these costs were very small, compared with the potential gains from trade. The minimum joint surplus two intramarginal traders could reap from a trade was 5 points (see Figure 1), i.e. 20 times the minimum costs of a bid.

Offers could be adjusted at any moment by simply submitting a new offer. Then the old offer 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, he was allowed to accept any offer of a potential trading partner of which he 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.

(b) No Transaction Costs (NTC) treatment

This treatment was identical to the TC treatment except that buyers (like sellers) did not have to bear transaction costs when making bids - bids were, like asks, for free. Hence, in this treatment the market environment was symmetric for buyers and sellers.

(c) Double Auction (DA) treatment

In the DA 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. There were no exogenous transaction costs. All traders were also informed when and at what price a trade occurred. This treatment boiled down to the standard continuous double-auction market, with the exception that trader IDs were added and that traders were not restricted to accept the best available offer.

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 the TC treatment the transaction 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. For treatments TC and NTC these lists only contained those offers that the trader was selected to be informed

about by the sending party. 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.

Procedure

For each of the three treatments, 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 hour decision making experiment which would earn them money. Fifteen subjects were registered for each experimental session to allow for no-shows. In session DA1, however, only 10 subjects showed up. This sessions was run with 10 traders, using the design of Figure 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 47, and ranged between Hfl 21 and Hfl 65. At the time of the experiment 1 Hfl exchanged for about \$ 0.50.

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 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 DA we have $D_{\text{client}} = 1$ and $D_{\text{comp}} = 1$ for all offers. All traders are informed about all asks and bids. In the TC and the NTC the two variables are endogenous and determined by the traders themselves.

In a previous study (Kirchsteiger, Niederle, and Potters 1998) we found that in absence of transaction costs offers tend to apply to all potential trading partners: $D_{\text{client}} = 1$. Competitors, however, were not informed: $D_{\text{comp}} = 0$. We referred to this endogenously emerging institution as a 'Secret Offer Double Auction' (SODA). Since asks were costless, there did not seem to be a reason why this pattern

should be different for asks in the TC sessions. However, due to the introduction of transaction costs we expected fewer bids than asks in the TC sessions - the majority of the offers should be asks in the TC treatment, whereas in the NTC treatment the number of asks and bids were equal. Furthermore, those bids actually made in the TC treatment should be disseminated less widely than the bids in the NTC treatment. Such an institution where most of the offers were asks disseminated among all buyers, but not among sellers, could be called "Secret Offer Auction" (SOA).

What did we expect regarding efficiency levels of the different treatments? As in numerous other DA experiments, we expected prices and quantities in our DA treatment to converge quickly to the market clearing level, and, in line with this, outcomes close to full efficiency. Quick convergence and high efficiency levels were also found in the SODA institution that emerged in the NTC treatment, although there traders typically observed only half of the offers and realized prices (for more details see Kirchsteiger, Niederle and Potters 1998).

Whether a market institution is (in)efficient in the presence of transaction costs is in general a subtle issue. On the one hand efficiency requires the transaction costs to be minimized. But a transaction cost minimizing institution might have disadvantageous information properties, since it might imply (too) little dissemination of the offers. This in turn could lead to slower convergence and to inefficiencies. Hence, overall payoff might be increased by incurring some transaction costs in order to allow the traders to reap larger gains of trade. Whether a certain institution provides the optimal trade-off between these two aspects is in general hard to say. Our design, however, allows for a solution to this problem. The endogenously emerging institution could be equivalent to an Offer Auction, where only sellers make asks and these asks are transmitted to all buyers and sellers. If such an institution emerges, no transaction costs have to be paid. Furthermore, previous studies showed that Offer Auctions lead to similar market outcomes as DAs (see Walker and Williams 1988). Hence, in the TC treatment an institution is feasible that i) implies no transaction costs, and ii) provides the same surplus to the traders as the DA. Therefore, if we find that earnings in the TC sessions were lower than earnings in the DA session, we can conclude that the emerging institution was inefficient.

If in the TC treatment a SOA institution resulted, then sellers received even less information about offers and realized prices than in the NTC treatment. Those who made the offers, namely the sellers, were only informed about their own offers and trades. This could have led to slower convergence and inefficiencies. Hence, we expected inefficiencies in case that a SOA actually emerged. This inefficiency, however, could have been a reason for a different institution than the SOA to arise.

3. Results

First we examine which market institution actually emerged endogenously in the TC treatment (Result 1)

and compare it with the institution emerging in an environment without transaction costs (NTC treatment). Then we compare efficiency levels (Result 2) across the three different treatments (TC, NTC, and DA).

Result 1. The Emerging Market is a SOA

The market institution emerging in the TC treatment is best described as a Secret Offer Auction (SOA). Most of the offers were made by the sellers, and these asks were typically sent to every buyer, but not to any competitor.

The numbers of asks and bids in the TC treatment provide a first evidence for Result 1. 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 1).

Insert Table 1 about here

In order to investigate the dissemination of the offers 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 maximum number to whom it could have been sent (i.e. divided by 6).³ 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.

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 NTC treatment the average D_{client} of bids was 0.858, it dropped in the TC treatment to a level of 0.384 (see Table 2a).

Insert Table 2 about here

On the other hand, the dissemination of asks did not change with the introduction of the transaction costs. Average D_{client} in the TC treatment was 0.853, whereas in the NTC it was 0.865 - typically asks applied to essentially all buyers (see Table 2a). 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 NTC to 3.3 in the TC treatment. On the other hand, the average number of asks a buyer gets per round did not change (14.3 in the NTC, 15.1 in TC treatment), and it was about 5 times as large as the number of bids a seller received in the TC treatment. Hence, we

³Note 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.

can conclude that the introduction of one-sided transaction costs induced the emergence of an institution where offers were made by the sellers, but not by the buyers.

Like in the NTC treatment, the dissemination of offers among competitors (D_{comp}) was very low in the TC treatment (see Table 2b). Hence, in both treatments offers were kept secret from the competitors.

Insert Figure 2 about here

Figure 2 shows that the average number of asks in the TC as well as the average number of asks and bids in the NTC treatment was between 15 and 20 during most of the rounds. Furthermore, these numbers were quite stable over the rounds. The average number of bids in the TC treatment, however, was considerably lower even in the very first round and declined even further during the course of the experiment. Hence, the more experienced the buyers the less they were willing to make bids in the TC treatment.

Figure 3 shows the evolution of the dissemination of offers (D_{client}) in the TC and the NTC treatment. From the very beginning nearly all offers in the NTC and nearly all asks in the TC treatment applied to all potential clients. The few bids made in the TC treatment, however, applied only to about 40% of the sellers, and no tendency towards a wider dissemination of bids was observed.

Insert Figure 3 about here

By and large, Result 1 is very favorable for the hypothesis that transaction costs are capable of shaping the market institutions. The introduction of one-sided transaction costs led to a sharp decrease of costly bids, and the few remaining bids were much less disseminated. But was the endogenously emerging institution also efficient? In order to answer this question, we have to distinguish between two different efficiency aspects of an institution. On the one hand, there is its ability to avoid transaction costs. Since the number of bids in the TC treatments was not zero, traders had to bear some transaction costs. But if the behavior of the buyers in the TC treatment would have been the same as it actually was in the NTC treatment, the transaction costs would have been more than 4.3 times as large as they actually were⁴ - the transaction cost induced change of the institution from SODA to SOA saved more than 75% of the transaction costs.

Insert Figure 4 about here

⁴This can be most easily seen by calculating the average transaction costs per round for the different treatments. They are given by: $(0.25) \times (\text{average number of bids}) \times (6 \times \text{average } D_{client} + 5 \times \text{average } D_{comp})$. In the TC treatment they were 5.1, i.e. 4.7% of the available surplus. Behavior as in the NTC treatment would have induced transaction costs of 22.7, i.e. 20.6% of the available surplus

Comparing the actual transaction costs with the hypothetical costs induced by behavior as in the NTC treatment, we observe that already in the very first round the differences in buyers' behavior were such that most of the transaction costs were saved (see Figure 4). Furthermore, there was a tendency for the already small costs to deteriorate further when the buyers became more experienced. Hence, the institution emerging in the TC treatment did not a perfect, but nonetheless a good job with respect to minimizing transaction costs.

The other important purpose of market institutions is the provision of a matching and information structure that allows the traders to reap the gains of trade. Hence, we had to investigate to what extent the endogenous institution in the TC treatment also served that purpose. We investigated the gains of trade not reaped by the traders, i.e. we looked to what the extent the sum of consumer and producer surplus fell short of its maximum level. Since here we analyzed the matching and information properties of the SOA institution, and not its ability to avoid transaction costs, the transaction costs in the TC treatment were not counted as forgone surplus. A natural measure of inefficiency is the foregone surplus as percentage of the maximum possible surplus. Another measure is the number of inefficient rounds, i.e. the number of rounds where the forgone surplus was larger than zero. Using either of these two measures we find:

Result 2. Efficiency

The forgone surplus was higher in the TC than in the NTC and the DA treatment.

Insert Table 3 about here

For each of the three treatments, Table 3 presents the average levels of foregone surplus and the number of inefficient rounds. It turns out that in the TC sessions the average inefficiency level was about 3%, whereas in the NTC and the DA sessions it was about 1%.⁵ The NTC sessions were about as efficient as the DA sessions. The efficiency levels of the individual sessions of the two treatments overlapped and the difference between the treatments was not significant ($p = 68\%$). This shows that an endogenous market institution can be as efficient as the DA. In case of transaction costs, however, we observe the emergence of a less efficient institution. Even the most efficient TC sessions (TC2 and TC4) were less efficient as the most inefficient NTC (NTC1) and DA session (DA4). Hence, the difference between the TC and the two other treatments was highly significant ($p = 2.8\%$)⁶. Qualitatively the same results appear if we look at the

⁵The impact of a trade of an extramarginal trader as well as that of a non-trade of an intramarginal trader on efficiency depends of course on the supply and demand conditions. Hence, not the absolute values of forgone surplus but the differences in efficiency between the treatments are important.

⁶Throughout we employed two-tailed Mann-Whitney tests with session averages as observations.

number of inefficient rounds. In the TC treatment 22 of 72 periods were inefficient, whereas in the NTC (DA) treatment only 9 (11) of 72 periods were inefficient. Again the difference between TC and NTC was significant ($p = 2.8\%$), whereas the difference between DA and NTC was not significant ($p > 50\%$).

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 non-trade 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 TC- than in the NTC and the DA treatment. We can conclude that the endogenously emerging SOA was more prone to mismatch (i.e. inefficient trades) as well as to non-match (i.e. inefficient non-trades) than the DA or the SODA.

Result 2 also shows that the endogenously emerging SOA was less efficient than the offer auction. In an offer auction only sellers make asks. The asks apply to all buyers and - unlike the SOA - 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 DA. Hence, if not the SOA but the offer auction would have emerged endogenously, foregone surplus would have been significantly lower. Furthermore, an offer auction causes of course no transaction costs in an environment like the one of our TC treatment. This implies that the sum of the net payoffs (taking into account transaction costs) would have been higher in the offer auction than it actually was in the SOA - the actually emerging institution is less efficient than another feasible one. Therefore, we cannot avoid the conclusion that the endogenously emerging SOA is inefficient, and also not the most efficient possible emerging market institution.

Why is the outcome of the TC market inefficient? We speculate that the SOA does not provide enough information for the 'active' side of the market, i.e. for the sellers. In a 'pure' SOA, when only asks are made and competitors are never informed, every seller knows only his own asks and only his own realized trades, whereas in an offer auction every offer and every acceptance is common knowledge. Of course, in the TC sessions the emerging SOA was 'impure', i.e. some bids were made and some asks were transmitted to other sellers. But the resulting information for the seller was too little to obtain an efficient outcome. In the NTC treatment an average seller was informed about 51.5 % of the offers made by others, and about 50.2% of the acceptance of offers made by others. In the TC treatment these numbers drop to 19.8% and to 10.8 %, respectively - on average a seller is only informed about every fifth offer and only about every tenth acceptance of an offer (other than his own offers). This dramatic loss in the information

provided to the market side making the offers seems to be the reason for the inefficiency. And if we look at the evolution of the institution over time, we observe no trend towards an institution providing more information, quite to the contrary.

Insert Figure 5 about here

While in the NTC treatment a seller was informed of about 50 % of all offers other than his own in all rounds, the average informedness of sellers in the TC treatment decreased from about 27 % in the first round to about 17 % in the last round (see Figure 5). This is mainly caused by the decline in the number of bids. The more buyers learn to avoid transaction costs, the less information sellers receive about the activities on the market. Of course, in a stable environment like the one of our experiment this problem is at least partly mitigated by the fact that sellers can accumulate information over the rounds. In an unstable environment, however, - like one with frequent changes of supply or demand - the problem of insufficient information for the active market side might become even more severe as subjects gain experience.

4. Concluding Remarks

In this paper we have investigated one potential driving force that determines the characteristics of market institutions, namely transaction costs. As postulated by New Institutional Economics our results have shown that even small transaction costs can shape the matching and information structure of markets. The resulting market institution, however, leads to inefficient trading outcomes. Hence, our results indicate that the widespread optimism concerning the self-organizing capacities of markets - if left alone, markets develop efficient institutional rules - is not warranted⁷. But of course this paper provides only one example for the emergence of inefficient institutions. To get a better understanding of the conditions for (in)efficient institutions to arise further investigations are needed. On the one hand one should test the robustness of the inefficiency result by conducting experiments with different exogenous transaction costs. If e.g. the structure of transaction costs are such that it becomes costly to hide information from the competitors it might be that market transparency increases and an efficient outcome can be restored. Varying transaction cost structure also allows for a test to what extent the emergence of institutions is driven by such transaction costs.

Furthermore, it is also important to endogenize the transaction costs. While one can think of some

⁷By investigating field experiments the same conclusion was also drawn by Roth (1984) and Roth and Xing (1994).

transaction costs that are caused by pure technological factors (and hence exogenous), most of the transaction costs seems to be caused by problems of incomplete information, incomplete contracts and imperfect contract enforcement. Hence, endogenous market experiments with moral hazard or adverse selection problems are a natural next step. But this is left to future research.

References

Abbink, Klaus, and Sadrieh, Abdolkarim (1995), RatImage - Research Assistance Toolbox for Computer-Aided Human Behavior Experiments, Discussion Paper No. B-325, University of Bonn.

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

Friedman, Daniel (1993), Double Auction Markets. A Survey, in: D. Friedman and J. Rust (eds.), *The Double Auction Market. Institutions, Theories, and Evidence*, Reading: Addison-Wesley.

Grether, David M., and Plott, Charles R. (1984), The Effects of Market Practices in Oligopolistic Markets: An Experimental Examination of the Ethyl Case, *Economic Inquiry* 22, 479-507.

Hayek, Friedrich. (1967), *Studies in Philosophy, Politics and Economics*, London: Routledge & Paul Kegan.

Holt, Charles A. (1995), Industrial Organization, in: J. Kagel and A. Roth (eds.), *The Handbook of Experimental Economics*, Princeton: Princeton University Press.

Hong James T., and Plott, Charles R. (1982), Rate Filing Policies for Inland Water Transportation: An Experimental Approach, *Bell Journal of Economics* 13, 1-19.

Kirchsteiger, Georg, Niederle, Muriel, and Potters, Jan (1998), The Endogenous Evolution of Market Institutions; CentER Discussion paper No. 9867.

Klein, B., Crawford, R., and Alchian, A. (1978), "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process", *Journal of Law and Economics* 21 (2), 297-326.

Lusht, Kenneth M. (1997), A Comparison of Prices Brought by English Auctions and Private Negotiations, *Real Estate Economics* 24 (4), 517-30.

North, Douglas C. (1990), *Institutions, Institutional Change and Economic Performance*, Cambridge: Cambridge University Press

North, Douglas C. (1993), Economic Performance Through Time, in: *The Nobel Prizes*, Stockholm: the Nobel Foundation.

Roth, Alvin E. (1984), The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory, *Journal of Political Economy* 92, 991-1016.

Roth, Alvin E., and Xing Xiaolin (1994), Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions, *American Economic Review* 84, 992-1044

Plott, Charles R. (1982), Industrial Organization Theory And Experimental Economics, *Journal of Economic Literature* 20, 1485-1527.

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

Smith, Vernon L. (1964), Effect of Market Organization on Competitive Equilibrium, *Quarterly Journal of Economics* 78, 181-201.

Stoll, Hans R., and Whaley, Robert E. (1990), Stock Market Structure and Volatility, *Review of Financial Studies* 3(1), 37-71.

Walker, James M., and Williams, Airlington W. (1988), Market Behavior in Bid, Offer, and Double Auctions, *Journal of Economic Behavior and Organization* 9, 301-314.

Williamson, Oliver. (1975), *Markets and Hierarchies: Analysis and Antitrust Implications*, New York: Free Press

Appendix: This appendix contains the instructions for the TC treatment. The instructions of the NTC treatment are exactly the same except for any reference to the transaction costs. For the DA treatment, 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 a "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. 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 hand-out you will only find the specific instructions that concern your role in the market.

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 asterix (*). 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 asterix (*). 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 $1 \text{ point} = 30 \text{ cents}$. 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. In earlier experiments with the present computer-program we have had some problems, especially with the network connections. Therefore, when entering something on your computer please wait a little bit until it is recorded on your screen before you enter something new. Also, 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.

Table 1. Average number of offers per round

Offer	Session	NTC	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 2a. Average values of D_{client} (all offers)

Offer	Session	NTC	TC
Bids	1	0.782	0.337
	2	0.875	0.411
	3	0.955	0.423
	4	0.820	0.366
	average	0.858	0.384
Asks	1	0.627	0.838
	2	0.965	0.810
	3	0.923	0.880
	4	0.945	0.884
	average	0.865	0.853

Table 2b. Average values of D_{comp}

Offer	Session	NTC	TC
Bids	1	0.099	0.009
	2	0.080	0.039
	3	0.091	0.016
	4	0.203	0.004
	average	0.118	0.017
Asks	1	0.023	0.117
	2	0.101	0.042
	3	0.090	0.074
	4	0.288	0.122
	average	0.125	0.089

Table 3. Foregone surplus and number of inefficient rounds

Session #	Foregone surplus (in % of total surplus)			number of inefficient rounds		
	DA	NTC	TC	DA	NTC	TC
1	1.48	1.77	3.54	3	3	5
2	0.25	0	2.02	1	0	5
3	1.26	1.26	3.79	3	4	7
4	1.77	0.76	2.02	4	2	5
All sessions	1.19	0.95	2.84	11	9	22

Note: In the TC treatment transaction costs did not count as forgone surplus, and did not count for the calculation of the number of inefficient rounds




Figure 1: The induced demand and supply schedule




Figure 2: Average number of offers




Figure 3: Average dissemination of asks and bids among potential trading partners (D_{client}) in the different treatments




Figure 4: Average actual transaction costs in the TC- and average hypothetical transaction costs in the NTC treatment




Figure 5: Percentage of other traders' offers on average known to a seller