



# Three minimal market institutions with human and algorithmic agents: Theory and experimental evidence <sup>☆</sup>

Juergen Huber<sup>a,b</sup>, Martin Shubik<sup>a</sup>, Shyam Sunder<sup>a,\*</sup>

<sup>a</sup> Yale University, 135 Prospect Street, New Haven, CT 06511, USA

<sup>b</sup> University of Innsbruck, Universitaetsstrasse 15, 6020 Innsbruck, Austria

## ARTICLE INFO

### Article history:

Received 8 July 2008

Available online 4 February 2010

### JEL classification:

C92

D43

D51

D58

L13

### Keywords:

Strategic market games

Laboratory experiments

Minimally intelligent agents

Adaptive learning agents

General equilibrium

## ABSTRACT

We define and examine three minimal market games (sell-all, buy-sell, and double auction) in the laboratory relative to the predictions of theory. These closed exchange economies have some cash to facilitate transactions, and include feedback. The experiment reveals that (1) the competitive general equilibrium (CGE) and non-cooperative (NCE) models are reasonable anchors to locate most but not all the observed outcomes of the three market mechanisms; (2) outcomes tend to get closer to CGE predictions as the number of players increases; (3) prices and allocations in double auctions deviate persistently from CGE predictions; (4) the outcome paths across the three market mechanisms differ significantly and persistently; (5) importance of market structures for outcomes is reinforced by algorithmic trader simulations; and (6) none of the three markets dominates the others across six measures of performance. Inclusion of some mechanism differences into theory may enhance our understanding of important aspects of markets.

© 2010 Shyam Sunder. Published by Elsevier Inc. All rights reserved.

## 1. Minimal market institutions

In this paper we define three minimal market institutions, and compare their theoretical properties to the outcomes observed in laboratory experiments with human agents and with simple algorithmic agents. These mechanisms are stripped of details and retain only the basic features necessary to be trading games playable in laboratory. Three price formation mechanisms considered here, listed by the nature of the strategy sets in a single market for each trader, are:

1. The sell-all model (strategy set of dimension 1);
2. The buy-sell model (strategy set of dimension 2)<sup>1</sup>;
3. The simultaneous double auction model (strategy set of dimension 2 or 4).

<sup>☆</sup> We thank Benjamin Felt and Ryan Dunn for their assistance with the laboratory experiments. We also thank the participants at University of Warsaw workshop and two anonymous referees and the editor for very helpful comments and suggestions that greatly improved the paper. Financial support by the Austrian Forschungsfoerderungsfonds (FWF, grant P-20609) is gratefully acknowledged.

\* Corresponding author at: Yale School of Management, 135 Prospect Street, New Haven, CT 06511, USA. Fax: +1 203 432 6974.

E-mail addresses: juergen.huber@uibk.ac.at (J. Huber), martin.shubik@yale.edu (M. Shubik), shyam.sunder@yale.edu (S. Sunder).

<sup>1</sup> Generically the dimensionality of the strategy set of the buy-sell model is two per market—the number of owned units of the good offered for sale and units of money bid to buy that good. In the laboratory implementation reported here, each individual was endowed with only one of the two goods, thus reducing the strategy set to dimension one per market—the number of owned units of one good offered for sale and the units of money bid to buy the other good.

These mechanisms utilize a commodity money for trade, and are described in Section 2. We find that non-cooperative and competitive general equilibrium solutions provide reasonable but imperfect static benchmarks to organize the laboratory observations. In absence of a widely accepted dynamic learning or disequilibrium theory, we compare the market outcomes of trading by profit-motivated humans to the outcomes of two simple computer simulations using minimally intelligent and adaptive learning algorithms as traders. The properties of even these minimal market mechanisms diverge when the number of traders is small. This differentiation raises questions of the appropriate level of specificity/generality for useful study of market mechanisms, to which we return in the final section of the paper.

The development of general competitive and non-cooperative equilibrium models has been followed during the recent decades by documentation of the properties of specific market institutions in game theory, industrial organization, experimental gaming, and experimental economics. The present study is an attempt to fill a gap that remained next to the abstract Walrasian end of the spectrum which is bereft of all institutional details.

Partial equilibrium exchange markets have been modeled as games in strategic form solved for their non-cooperative equilibria starting with Cournot (1838, 1897), Bertrand (1883), and Edgeworth (1925), followed by many others. Nash (1951), provided the full generalization of the concept of a non-cooperative equilibrium and Dubey (1982), Dubey and Shubik (1978, 1980), Quint and Shubik (2005, 2007), Shapley (1995), Shapley and Shubik (1977), Shubik (1973), Sorin (1996) and several others extended the analysis to closed economies. There is also a related partial equilibrium literature introducing uncertainty into auction and double auction models as is evinced by the work of Vickrey (1961), Griesmer et al. (1967), Milgrom and Weber (1982), and Satterthwaite and Williams (1989).

There are two other relevant literatures: one in macro-economics stressing rational expectations (exemplified by Lucas, 1987, 1988; Lucas and Sargent, 1981) and the other in mathematical finance mostly on competitive partial equilibrium open models dealing explicitly with money, transactions costs, and the constraints on cash flows. All approaches broadly involve money, markets and financial institutions. There has been considerable gaming activity on bargaining, bidding and on the emergence of competitive prices in some simple markets with little stress on the explicit role of money (Marimon et al., 1993; Lim et al., 1994; Marimon and Sunder, 1993, 1994, 1995). Our paper presents gaming with a role for money; three other papers include credit and other financial instruments in addition to money (Angerer et al., 2009; Huber et al., 2008, 2009).

Experiments that examine the properties of markets and competition (Smith, 1982; Plott, 1982) show that markets with only a few independent individual traders often yield outcomes in close neighborhood of competitive equilibrium predictions. Most experimentation has involved trade in a single market. In the spirit of general equilibrium, we consider two markets. We formulate experimentally playable strategic market games where the trade is mediated by money, but the overall system is closed. Unlike open or partial equilibrium settings of most other experiments, these closed exchange economies have limited amounts of cash to facilitate transactions, and include feedback. The experiment reveals that (1) the competitive general equilibrium (CGE) and non-cooperative (NCE) models provide a reasonable anchor to locate most but not all the observed outcomes of the three market mechanisms; (2) outcomes tend to get closer to CGE predictions as the number of players increases; (3) prices and allocations in our double auctions with full feedback reveal significant and persistent deviations from CGE predictions; (4) the outcome paths from the three market mechanisms exhibit significant and persistent differences among them; (5) since the dynamics of markets populated by profit-motivated human subjects is at least partially captured in markets with simple algorithmic traders, the importance of market structures in determining their outcomes is reinforced; and (6) none of the three markets necessarily dominates the others across a set of six measures of performance. These are allocative efficiency, money balances, symmetry of investments, price levels, the dispersion of profits, and the trading volume in the markets. The results suggest that abstracting away from all institutional details does not help understand dynamic aspects of market behavior and that inclusion of mechanism differences into theory may enhance our understanding of important aspects of markets and money, and help link conventional analysis with dynamics.

The remainder of the paper is structured as follows: In Section 2 the three minimal institutions are described. Section 3 gives the general and non-cooperative equilibrium predictions for each institution which serve as static benchmarks for comparing the experimental data. Section 4 describes two dynamic benchmarks—minimally intelligent<sup>2</sup> and adaptive learning algorithmic traders. Section 5 describes the experimental setup we used to implement these markets in the laboratory with human traders. The results are presented in Section 6, followed by some concluding remarks.

## 2. Three minimal market games

We examine three mechanisms which are minimal in the following sense. In order to reflect an exchange economy with money we need at least two commodities in addition to money whose special properties we wish to explore. A game cannot have less than one information set and less than one move per player. If they move simultaneously there will be one information set. Further, price should be at least generically sensitive to, i.e., be a function of, bids and offers. In the sell-all game, the money bid for each commodity is the single move in each market, and calculation of price as the ratio of the sum of money bid and total available quantity of the commodity is the simplest price function. If the mechanism is to

<sup>2</sup> Since Gode and Sunder's "zero intelligence" agents originally defined for double auctions had to be modified to operate in broader classes of market environments, we changed the label to "minimally intelligent."

satisfy an additional requirement that agents either buy or sell (and possibly do both) in the market for each commodity, we get the buy–sell as the minimal mechanism; the strategy set still has dimension 1 although it consists of the quantity of endowed good offered for sale in one market, and the quantity of money bid in the other market. Finally, the requirement that individuals be able to specify their price and quantity limits leads to a double sealed bid as the minimal mechanism with a four dimensional strategy set, although we use sequential double auction in this paper because its properties have been studied extensively in the experimental gaming literature. It differs from the double auction sealed bid in the number of information sets.

## 2.1. Definitions

### 2.1.1. Money

In each market game two commodities are traded and one more instrument is used as a means of payment (money). This money is introduced as a linear term in the subjects' utility functions.<sup>3</sup>

### 2.1.2. Bids

**2.1.2.1. A money bid** A trader  $i$  bids an amount of money  $b_j^i$  for the  $j$ th commodity. The trader has no reserve price and accepts the market price. This allows a simple quantity bid for a mechanism similar to Cournot's 1897. The market clearing mechanism gives the trader  $i$  quantity  $x_j^i = b_j^i/p_j$  of good  $j$  where  $p_j$  is the market price that is formed collectively from individual bids and offers.

**2.1.2.2. A price-quantity bid** Suppose that a trader  $i$  instead of offering an amount of money to buy a good  $j$ , bids a personal unit price  $p_j^i$  he is willing to pay to buy up to an amount  $q_j^i$  of the good. It is reasonable to expect that he is willing to buy  $q_j^i$  or less at a price less than or equal to  $p_j^i$ . There is an implicit limit in this bid inasmuch as  $q_j^i p_j^i$  must be less than or equal to the individual's credit line plus cash. Since we do not consider a credit mechanism in the three market institutions considered here,  $q_j^i p_j^i$  cannot exceed the available cash. Minor variations of these bids consider any upper or lower bounds on prices or quantities acceptable to the bidder.

### 2.1.3. Offers

Analogously, there are two simple forms of offers.

**2.1.3.1. A non-contingent offer to sell** Suppose that an individual  $i$  owns  $a_j^i$  units of good  $j$  and wishes to sell some of it. The simplest strategy is for her to offer  $q_j^i \leq a_j^i$  units for sale at the market-determined price.

A somewhat more complex action, but still not involving any more information and confined to a single move is:

**2.1.3.2. The price-quantity offer** Suppose that a trader  $i$  is willing to sell up to an amount  $q_j^i (\leq a_j^i)$  of good  $j$  at unit price  $p_j^i$ . It is reasonable to expect that she is willing to sell  $q_j^i$  or less for a price more than or equal to  $p_j^i$ .

We use observable acts to buy (bids) and sell (offers) as the building blocks to construct three simple market games. Simplifying them any further will prevent any trading. The first two market games involve a single move by every agent, taken simultaneously. The third, double auction, involves sequential multiple moves by various players. Each game can be generalized to multiple plays.

Consider  $n$  individuals where  $i$  has an endowment  $a_j^i$  of good  $j$  ( $j = 1, \dots, m$ ) and an endowment  $M^i$  of money. Suppose there are  $m$  markets, one for each good  $j$  where it can be exchanged for money. A plausible restriction on the market mechanism is that all trades in a given market take place at the same time and the same price. This requires that  $p_j^i = p_j$  for  $i = 1, \dots, n$ .

In general, we cannot assume that bids in one market are independent of bids in the others. There is at least a cash or credit budget constraint that links actions across markets.

## 2.2. The sell-all model

This is the simplest of the three models. Consider  $n$  traders trading in  $m + 1$  goods, where the  $m + 1$ st good has a special operational role, in addition to its possible utility in consumption. Each trader  $i$  has an initial bundle of goods and money  $a^i = (a_1^i, \dots, a_m^i, M^i)$ , where  $a_j^i \geq 0$  for all  $j = 1, \dots, m + 1$  and  $a_{m+1}^i = M^i$ , and the utility  $u^i = u^i(x_1, \dots, x_m, x_{m+1})$ , where  $u^i$  need not actually depend directly on  $x_{m+1}$ ; a fiat money is not excluded.

In order to keep strategies simple, let us suppose that the traders are required to offer for sale *all* of their holdings of the first  $m$  goods. Instead of owning their initial bundle of endowments outright; the traders own a *claim on the proceeds* when the bundle is sold at the prevailing market price.

<sup>3</sup> For a detailed justification for this assumption see Shubik (1999). Many properties are attributed to a money, but the central one studied here is the "means of payment."

Suppose there is one trading post for each of the first  $m$  commodities, where the total supplies  $(a_1, \dots, a_m)$  are deposited for sale “on consignment,” so to speak. Each trader  $i$  submits bids by allocating amounts  $b_j^i$  of his endowment  $M^i$  of the  $m + 1$ st commodity among the  $m$  trading posts,  $j = 1, \dots, m$ . There are a number of possible rules governing the permitted range of bids. In the simplest case, with no credit of any kind, the limits on  $b_j^i$  are given by:

$$\sum_{j=1}^m b_j^i \leq M^i \quad \text{and} \quad b_j^i \geq 0, \quad j = 1, \dots, m.$$

An interpretation of this spending limit is that the traders are required to pay *cash in advance* for their purchases. The prices are formed from the simultaneously submitted bids of all buyers; we define

$$p_j = b_j / a_j, \quad j = 1, \dots, m.$$

Thus, bids precede prices. Traders allocate their budgets fiscally, committing specific quantities of their means of payment to the purchase of each good without definite knowledge of what the per-unit price will be (and how many units of each good their bid will get them). At an equilibrium this will not matter, as prices will be what the traders expect them to be. In a multi-period context, moreover, the traders will know the previous prices and may expect that variations in individual behavior in a mass market will not change prices by much. But any deviation from expectations will result in changing the quantities of goods received, and not in the quantities of cash spent. In a mass market, the difference between the outcomes from allocating a portion of one's budget for purchase of a certain good, and from a decision to buy a specific amount at an unspecified price, will not be too different.

The prices in the model are determined so that they exactly balance the books at each trading post. The amount  $x_j^i$  of the  $j$ th good that the  $i$ th trader receives in return for his bid  $b_j^i$  is

$$x_j^i = \begin{cases} b_j^i / p_j & \text{for } p_j > 0, \quad j = 1, \dots, m, \\ 0 & \text{if } p_j = 0, \quad j = 1, \dots, m. \end{cases}$$

His final balance of the  $m + 1$ st good, taking account of his sales as well as his purchases, is

$$x_{m+1}^i = a_{m+1}^i - \sum_{j=1}^m b_j^i + \sum_{j=1}^m a_j^i p_j.$$

### 2.3. The buy-sell model

Subjects face a more complex task in the buy-sell model: instead of one money bid in each of the two markets in sell-all, they submit the quantity  $q_j^i$  of their endowed good they wish to sell, and the money bid for the other good they want to buy. Thus they enter only one number in each market but these numbers are in different dimensions (goods and money). Since moves are simultaneous, there are no contingencies in this market either. Physical quantities of goods are submitted for sale and quantities of money are submitted for purchases, and the markets are cleared. The mechanism does not permit any underemployment of resources.<sup>4</sup> The amount  $x_j^i$  of the  $j$ th good that the  $i$ th trader receives in return for his bid  $b_j^i$  is:

$$x_j^i = \begin{cases} b_j^i / p_j & \text{if } p_j > 0, \quad j = 1, \dots, m, \\ 0 & \text{if } p_j = 0, \quad j = 1, \dots, m. \end{cases}$$

However price is somewhat different as it depends on the quantities of each good offered for sale (and not on the endowment of each good):

$$p_j = b_j / q_j, \quad j = 1, \dots, m.$$

His final amount of the  $m + 1$ st good, taking account of trader  $i$ 's sales as well as his purchases, is

$$x_{m+1}^i = a_{m+1}^i - \sum_{j=1}^m b_j^i + \sum_{j=1}^m q_j^i p_j.$$

### 2.4. The sequential bid-offer or double auction model

Any trader is free to submit a bid in either market to buy one unit at or below a specified price, and an ask to sell one unit at or above a specified price as long as he has the money (to buy) or the good (to sell). The computer screen shows all outstanding bids in descending order and all outstanding asks in ascending order. Traders are free to accept the lowest

<sup>4</sup> Except when there is no bid or offer, in which instance all resources are returned to their owners. If they are ripe tomatoes, the owner is in trouble.

**Table 1**Non-cooperative equilibria for sell-all model (parameter values used in the laboratory experiments:  $a = 200$ ;  $M = 6000$ ;  $\alpha = 10$ ).

Number of agents	Price (1)	Price (2)	Quantity of endowed good	Quantity of other good	Unspent money	Payoff
2	21.14	21.14	0.6277a	0.3723a	0.2953M	0.4834 $\alpha a$
3	20.40	20.40	0.5838a	0.4162a	0.3200M	0.4929 $\alpha a$
4	20.20	20.20	0.5626a	0.4374a	0.3267M	0.4961 $\alpha a$
5	20.12	20.12	0.5501a	0.4499a	0.3293M	0.4975 $\alpha a$
10	20.03	20.03	0.5250a	0.4750a	0.3323M	0.4994 $\alpha a$
Many	20.00	20.00	0.5000a	0.5000a	0.3333M	0.5000 $\alpha a$

outstanding bid or the highest outstanding ask and consummate a trade. If the highest bid and lowest ask cross, a trade is automatically recorded at that price.

The double auction model doubles the size of the strategy set, changing price into a strategic variable from a mere outcome of the quantity strategies in the sell-all and buy-sell models. In each of the  $m$  markets, an individual's strategy has four components  $(p, q; p^*, q^*)$  where the first pair of numbers is interpreted as an offer to sell amount  $q$  or less for a price  $p$  or more, and the next pair is a bid to buy amount  $q^*$  or less at a price  $p^*$  or less.

From the viewpoint of both game theory and experimental gaming the number of decisions in a double auction is more than in the other two markets. Imposing a condition that one can either buy or sell, but not both, is a possible theoretical simplification. In practice, however, an individual can be a buyer or a seller or a trader. Most consumers are buyers and most producers are sellers of specific commodities or services; a trader can be active on both sides of the market.

In these games the terminal amount of money  $(M - b + pa)$  held by each individual was added to their dollar payoffs. This served to fix the price level that the transactions would be expected to approach towards the end. The observed divergence between these predicted and realized prices in some cases was considerable, and is discussed later.

### 3. General and non-cooperative equilibria

The non-cooperative equilibrium (NCE) solution is a fairly natural game theoretic way to approach these games without any direct communication. A non-cooperative equilibrium satisfies the existence of mutually consistent expectations. If each predicts that the other will play his strategy associated with a non-cooperative equilibrium the actions of all will be self-confirming. No one acting individually will have an incentive to deviate from this equilibrium. This could be called an outcome consistent with "rational expectations," but as the outcome may be neither unique nor generically optimal, the label of "rational" is best avoided.

The competitive general equilibrium (CGE) solution is defined as the set of prices that clear all markets efficiently. In general, the mathematical structure of NCE and CGE differ. However, it can be shown in theory that, as the number of traders in a market increases, under reasonable conditions, the NCE approaches the CGE. In symmetric markets without face-to-face communication experimentation can verify that with as few as 5–10 traders on each side, the outcomes approximate the CGE, and any differences between the two can be explained by the NCE.

#### 3.1. The non-cooperative equilibrium in the sell-all market

Sell-all is the simplest model and for experimental purposes we keep the payoff structure simple to explain to subjects untutored in economics or mathematics:

$$\alpha\sqrt{xy} + M - b + pa,$$

where  $\alpha$  is an appropriately chosen parameter (explained in the discussion of the game), the square root of  $xy$  is a simple Cobb–Douglas utility function whose range of values is furnished in a coarse-grid table in order to ease the computational burden. The linear term  $(M - b + pa)$  is the residual amount of money (initial endowment less the amount of money bid plus earnings from selling  $a$  units at price  $p$ ).<sup>5</sup>

The mathematical solutions of this model under different constraints are given in Appendix A. Table 1 shows the NCE for markets with 2, 3, 4, 5, 10 and many traders on each side for the parameter values used in the experiment.

#### 3.2. The non-cooperative equilibrium in the buy-sell market

The basic difference between the sell-all and the buy-sell model lies in the freedom subjects have to control the amount of goods to sell in the market for the endowed good (see Table 2). The general formulae for the NCE are given in Appendix A.

<sup>5</sup> The utilization of a money with a Marshallian or constant marginal utility can be interpreted in terms of a known expectation of the worth of future purchasing power. In this context any change in price level can be attributed to error and learning the equilibrium of the actual game is stationary. This device provides an easy and logically consistent way in an experimental game to provide terminal conditions.

**Table 2**Non-cooperative equilibria in buy–sell market (parameter values used in the laboratory experiments:  $a = 200$ ;  $M = 4000$ ;  $\alpha = 10$ ).

Number of agents	Price (1)	Price (2)	Quantity of endowed good	Quantity of other good	Unspent money	Payoff
2	20.00	20.00	0.8000a	0.2000a	0.8000M	0.4000 $\alpha a$
3	20.00	20.00	0.6923a	0.3077a	0.6923M	0.4615 $\alpha a$
4	20.00	20.00	0.6400a	0.3600a	0.6400M	0.4800 $\alpha a$
5	20.00	20.00	0.6098a	0.3902a	0.6098M	0.4878 $\alpha a$
10	20.00	20.00	0.5525a	0.4475a	0.5525M	0.4972 $\alpha a$
Many	20.00	20.00	0.5000a	0.5000a	0.5000M	0.5000 $\alpha a$

### 3.3. The non-cooperative equilibrium in the double-auction model

For simplicity, the bid-offer market is modeled as a simultaneous sealed bid auction. The clearing method for the one-shot game is simplicity itself. Bids are assembled in a down-sloping histogram and offers in an up-sloping histogram. Market price is formed where the two lines intersect.<sup>6</sup>

The double auction used in the experiment is a continuous process where bids and offers flow in sequentially and a trade takes place whenever they match or cross. We use this continuous double auction rather than the simultaneous sealed bid auction so traders can learn from the order-book and from past prices.

Two individuals on each side of the market are sufficient for the competitive equilibrium to be a NCE. A simple example considering optimal response is sufficient to show this. Suppose that there are two individuals each of two types. All have the payoff function given above, but individuals of type 1 and 2 have endowments of  $(a, 0, M)$  and  $(0, a, M)$ , respectively, where the first component is the endowment of the first good, the second the endowment of the second good and the third the endowment of money. Suppose  $M > a/2$  and  $\alpha = 2$  (the parameter in the payoff function), a trader of type 1 offers to sell  $a/2$  or less of good 1 at a price of 1 or more and to buy up to  $a/2$  of good 2 at a price of 1 or less, it can be verified that this is an equilibrium outcome and the price of both goods is 1 ( $p_1 = p_2 = 1$ ).<sup>7</sup>

There is a considerable amount of experimental evidence that in a single market the double auction mechanism yields efficient allocations. In their single-commodity double auctions, Gode and Sunder (1993, 1997) found that it requires negligible skills or intelligence from traders for the market outcome to lie in close proximity of the competitive equilibrium. However, we consider two markets for two commodities; whether the complementarities between the two make a difference remains open. Obviously the task of trading on two markets simultaneously is markedly more demanding than trading on a single-commodity market.

In their one-shot versions, the three games are the simplest price formation mechanisms that can be constructed, involving the maximum of one (sell-all and buy–sell) and four (double auction) strategic variables. They can all be analyzed for their NCE. Unlike most other market experiments, these are general equilibrium full feedback models, not partial equilibrium constructs.

The non-cooperative model of the general equilibrium in theory, generates an asymmetry in actions when there are few agents, as can be seen in the sell-all model where a seller obtains an oligopolistic income from buying a commodity to which he has ownership claims (as contrasted with buying a commodity he does not have). This asymmetry is the largest in the buy–sell game, the next largest in the sell-all game and the smallest in the double auction (see Tables 1 and 2 for numerical examples for 5 + 5 traders).

Paradoxically, because minimally intelligent (MI) agents (see Section 4 below) ignore their oligopolistic influence the theoretical prediction is that in all markets the price should be as close or closer to the competitive equilibrium than with oligopolistic human traders, but because of the random action there should be a variation in payoffs that is not present in the equilibrium analysis of the three games.

The speed of learning and the variation among players is not predicted by the static non-cooperative or general equilibrium theories. Many learning theories have been proposed and in the next section, we consider one non-learning and one simple learning algorithm. We only conjecture that as human subjects learn, variations in the outcomes of markets will diminish in the later periods (replications) of the game.

## 4. Dynamic benchmarks

Richness of the data sets generated from market experiments with human subjects is not captured in the static NCE and CGE benchmarks. Unfortunately, there is no generally accepted disequilibrium theory of dynamic learning. We therefore compare the results obtained from markets populated by profit-motivated human traders with the results from markets populated by two different kinds of simple algorithmic traders described in this section: the non-learning minimal in-

<sup>6</sup> It is necessary to take care of several cases; see Dubey and Shubik (1980) or Dubey (1982).

<sup>7</sup> From a strictly technical game theoretic point of view there is a continuum of non-cooperative equilibria, all with the same efficiency that are consistent with the competitive equilibrium outcome.

telligence (MI) benchmark (after Gode and Sunder, 1993's zero-intelligence or ZI, see footnote 2), and adaptive learning agents (AL).

#### 4.1. Minimally Intelligent (MI) traders

In sell-all markets, given the money endowment of  $M$ , each agent picks a uniformly distributed random number between 0 and  $M$  as its total money bid (for goods A and B combined). A second uniformly distributed random variable  $z$  between 0 and 1 is drawn to define the share of this money bid invested in A with  $(1 - z)$  invested in B.

In the buy–sell market, each trader offers to sell a randomly chosen quantity of the endowed good (from uniform distribution between 0 and  $a$ ) and bids a randomly chosen quantity of money for the other good (from uniform distribution between 0 and  $M$ ).

In double auctions, with equal probability and independently, one trader is picked, one of the two markets is picked, and either bid or ask is picked. Given the trader's current holdings of the two goods and cash, the computer calculates the opportunity set (the maximum amount of bid the agent can make without diminishing its net payoff), and draws a random number between the current bid and this calculated upper limit (if the latter is more than the former) and submits it as a bid from this trader. In case of asks, the computer calculates the minimum amount of ask the trader can submit without diminishing its net payoff and submits a random number between this calculated lower limit and the current ask (if the latter is above the former), as the ask.<sup>8</sup> Higher bids replace lower ones as market bids, and lower asks replaced higher ones as market asks. Whenever market bids and market asks cross, a transaction is recorded at the price equal to the bid or ask, depending on which of the two was submitted earlier (see Appendix B).

#### 4.2. Adaptive Learning (AL) traders

The adaptive learning (AL) algorithmic traders are a modification of the MI traders described in the preceding paragraphs. In sell-all and buy–sell markets, each AL trader keeps track of the past decisions which yielded the highest payoff and uses an adaptive learning parameter  $\lambda$  (set to 0.5 in the simulations) to adjust the most recent decision towards this “historical best” decision. The bid for the next period is then  $\lambda$  times the “historical best” decision plus  $(1 - \lambda)$  times new random variables (as in MI).<sup>9</sup>

In double auction algorithm starts period 1 with a “price aspiration” of money/goods in the endowed quantities and uses each observed transaction price to adjust this aspiration by  $\lambda^*(\text{transaction price} - \text{price aspiration})$ . In addition to the constraints described above in the description of MI traders, AL traders use this price aspiration as an additional constraint, not bidding above and not asking below this level. We chose learning algorithms where the agents only look at their own earnings and their own decisions; market variables are not considered.

The paths of markets populated by these two kinds of artificial players should serve as much as a warning as benchmarks. Rigid rule gaming in cleaned up abstract laboratory conditions contrasts sharply with the battlefield conditions of phenomena of substantive interest. Under the conditions chosen here, there is a unique analytical interior perfect non-cooperative equilibrium. In such situations, it is not difficult to find many dynamic procedures such as hill-climbing, optimal response, exponential lag weighted forecasting or adaptive forecasting rules that work well on a reasonably smooth terrain with a unique joint maximum. Kumar and Shubik (2004) note that one can take an example such as the well-known Scarf model of global instability with a unique equilibrium point and easily find a control process that gives contrary results.

The large body of work that applies dynamic programming microeconomic methods to problems of macroeconomics tells us little about learning and disequilibrium behavior. Our human and algorithmic games merely yield an empirical picture of the markets populated by various kinds of traders. It is easy to fit many process models to the data *ex post*. We, too, could try to fit some plausible rules of behavior to the observed data. The gains from such an exercise being doubtful, we refrain from doing so.

### 5. The experimental setup

We conducted and report on two independent sessions for each of the three market games considered in this paper. In each session, programmed in Z-tree software (see Fischbacher, 2007), the participants traded two goods—labeled A and B—for money. Each session had ten participants, five of them were endowed with some units of A and none of B, while the other five had some units of B and none of A.<sup>10</sup> All had the same starting endowment of money. Each session consisted of ten or twenty independent rounds of trading. Subjects' “consumption” at the end of each round was accumulated in a “bank account” with the experimenter. No goods balances were carried over from one round to the next, and each subject was re-endowed with the ownership claims to goods A or B at the beginning of each round. In all treatments money is carried over to the following round (see descriptions of specific treatments below and in Table 3).

<sup>8</sup> This means that bids are randomly distributed  $\sim U(\text{Current Bid}, (100/0.5)((c_A + 1)c_B)^{0.5} - (c_A c_B)^{0.5})$ ; asks are randomly distributed  $\sim U((100/0.5) \times (-(c_B - 1)c_A)^{0.5} + (c_A c_B)^{0.5}, \text{Current Ask})$ . After each transaction, current bid is set to 0 and current ask is set to the initial cash balance of 4000.

<sup>9</sup> With  $\lambda = 0$  the AL-simulation would be the same as the MI-simulation as then no learning would take place.

<sup>10</sup> In addition, we conducted one session with 20 participants, of whom 10 were of each type.

**Table 3**

Design parameters for six sessions of three market games.

Runs	Market game	Endowments of individuals			Money carried over?	Payoff function
		Good A	Good B	Money		
1 + 2	Sell-all	200 for 5 traders; 0 for 5 others	0 for 5 traders; 200 for 5 others	6000	Yes	$10(c_A c_B)^{0.5}$ each period + 0.25 final money bal
3 + 4	Buy–sell	200 for 5 traders; 0 for 5 others	0 for 5 traders; 200 for 5 others	4000	Yes	$10(c_A c_B)^{0.5}$ each period + 0.5 final money bal
5 + 6	Double auction	20 for 5 traders; 0 for 5 others	0 for 5 traders; 20 for 5 others	4000	Yes	$100(c_A c_B)^{0.5}$ + 0.5 final money bal

### 5.1. Sell-all call market

In Treatment 1 (sell-all market) the initial endowments were 200 or 0 units of A, 0 or 200 units of B, and 6000 in cash. All units of A and B were sold automatically at a price derived from the set of bids submitted by the traders. In other words, subjects did not have to decide on the number of units they wished to sell; all their holdings of goods were sold at the prevailing market price. Consequently, they had ownership claim to the revenue from selling 200 units of the good they were endowed with. The only decision participants had to make was how much of their money endowment they wished to bid to buy good A and how much to bid to buy good B. Each sell-all market was repeated for 20 periods.<sup>11</sup>

As outlined above the unit prices of A and B are calculated as the respective sums of money bid for the respective good by all traders divided by the total units of each goods for sale. With 6000 units of money endowment per trader there is more than enough money to reach general equilibrium at prices of 20 per unit of A and B. At general equilibrium traders would spend 2000 on each good and keep 2000 of their money endowment unspent. However, in a thin market with only a few traders, deviating from general equilibrium spending level may make sense to traders. When a trader spends more on the good he is endowed with, he raises its price and therefore his revenue from selling a part of his endowment of this good. Apart from the general equilibrium, there also exists a non-cooperative equilibrium in which traders spend 2213.4 on the good they own, 1810.6 on the other good, and keep 1976.0 unspent. Prices are slightly higher at 20.12 for both goods in this equilibrium. We conducted two runs of this treatment.

### 5.2. Buy–sell call market

Unlike in Treatment 1, traders in this treatment directly control the goods they are endowed with, and decide how many, if any, units they wish to sell (in Treatment 1 all units were sold automatically). Again half of the traders are endowed with 200 units of A and none of B, while the other half are endowed with 200 units of B and none of A. Each trader has an initial endowment of 4000 units of money at the beginning of the first round of the session. Money balances are carried over from one round to the next. Each buy–sell market was repeated for 20 periods.

Traders make two decisions: The amount of their money to buy the good they do not own, and the number of units to sell out of the 200 units of the good they own.

Prices for A and B are calculated by dividing the total investment for the respective good by the number of units put up for sale. Competitive equilibrium prices and conditions are the same as in Treatment 1. Final holdings of goods are (100, 100) each (prices are 20/20, each trader spends 2000 for the good he does not own, and sells 100 units of the good he owns). At the non-cooperative equilibrium with 5 traders on each side of the market traders of type 1 offer 78.05 units of the second good for sale and bid 1560.97 units of money for the first good. Traders of type 2 do the opposite (see Table 2). Final consumptions are (78.05, 121.95) for traders of type 1 and (121.95, 78.05) for traders of type 2. Prices are 20/20.

### 5.3. Double auction market

Treatment 3 features a double auction market where participants can trade goods A and B in a continuous market for several periods. We simplify trading by considering only transactions for one unit at a time. To reduce the number of transactions needed to reach equilibrium levels, initial endowments of A and B are reduced to (20/0, 0/20), so traders own 20 units of a good rather than 200. Each period lasts for 180 seconds to give participants enough time to buy ten units of goods they do not yet own, and sell ten units of the good they are endowed with, required to reach equilibrium. Traders are endowed with 4000 units of money, which is more than enough for trading.

<sup>11</sup> Instructions and the trading screens for this as well as other treatments are available as supplementary material to this article.



Competitive equilibrium and non-cooperative equilibrium prices coincide for the closed double auction model as was shown by Dubey (1982).<sup>12</sup> They are 100 for each good. The first run of the double auction market was repeated for 10 periods, the second run for 11 periods.

In the double auction sessions we allow market as well as limit orders. All orders are executed according to price and then time priority. Market orders have priority over limit orders in the order book. This means market orders are always executed instantaneously. Again, holdings of money are carried over from one round to the next, while holdings of goods are reinitialized.

Participants receive current information about their cash and stock holdings, their wealth, and their transactions within the current period on the screen. In the centre of the screen they see the open order books and they have the opportunity to post limit or market orders. On the left side of the screen transaction prices of the round are charted against time.

## 6. Results from the experiments

In Tables 5 and 6, and in various panels of Figs. 1–8, we use six aspects of market outcomes—allocative efficiency, prices, symmetry of allocation across the two goods, money balances (except in double auctions where it is undefined), cross-trader dispersion of earnings, and trading volume—to assess the behavior of three market mechanisms relative to three static (autarky, competitive general equilibrium, and non-cooperative equilibrium), and two dynamic (markets populated by minimally intelligent or MI, and adaptive learning or AL agents) benchmarks.

Allocative efficiency of the markets is measured each period by the average amount earned by traders as a percentage of the competitive general equilibrium amount (1000 points = 100%). Six panels of Fig. 1 show the time series of efficiency in two replications of each of the three market games; in the left column of panels the human market data are charted against the background of quintiles of efficiency statistics from 1000 replications of markets populated by MI algorithmic traders and the right column has quintiles from AL algorithmic traders in the background. The autarky (efficiency = 0) benchmark is not included in the chart. The solid black line of competitive general equilibrium (efficiency = 100) frames the charts at the top and the non-cooperative equilibrium efficiency (for  $5 + 5 = 10$  players) is shown in a dotted line slightly below.

The development of transaction prices is measured by market clearing prices for sell-all and buy-sell markets, and by average transaction prices (averaged across transactions within one period) in the double auction markets. Each of the six panels in Fig. 2 charts the observed prices (or average prices in DA) for goods A and B in two sessions of one of the three market games. The continuous horizontal line marks the CGE prices (20 in sell-all and buy-sell and 100 in DA). The corresponding NCE prices are 20.12, 20, and 100, and therefore do not show as a separate line in these charts. The data are charted against the background of gray quintile bands of prices from MI simulations in the left column; the same data are charted again in the right column of panels against the background of quintiles from AL simulations.

Symmetry of allocation is the ratio of consumption of good A and B ( $= \min(c_A/c_B, c_B/c_A)$ ). Given the parameters chosen for these experiments, goods A and B should be allocated symmetrically at the competitive equilibrium, which has the symmetry measure of 1. Autarkic symmetry is 0.

Money balances refer to the percentage of initial money left unspent after buying decisions are made (and before the proceeds of any sales are received) in sell-all and buy-sell markets.

We report these four performance measures relative to the three above mentioned static benchmarks summarized in Table 4. Under autarky, efficiency and symmetry are 0, prices are undefined, and money balance is 100 percent. The competitive general equilibrium allocations are 100 units each of good A and B in sell-all and buy-sell markets, and 10 units of each good in the double auctions, yielding a symmetry measure of 1 in all cases. Prices are 20 in sell-all and buy-sell markets, and 100 in the double auction markets.

The third benchmark for market performance is non-cooperative equilibrium for 10 traders (five endowed with good A and five endowed with good B). Application of theory to the parameters of these markets yields bids of 2214 and 1811 for the owned and the not-owned good, respectively, and final holdings of 110 and 90 units in the sell-all model. In buy-sell model non-cooperative equilibrium requires selling 78 of the 200 units of the owned good and buying 78 with a bid of 1561 units of money. In the double auction traders should keep 11 of their 20 units of the good they are endowed with and buy 9 of the other. Unspent money balance is 32.92 percent in sell-all, 60.98 percent unspent in buy-sell, and not defined in the double auction. The resulting measures for symmetry are 0.82 in sell-all, 0.64 in buy-sell, and 0.82 in double auction. Prices are 20.12 in sell-all, 20 in buy-sell, and 100 in double auction.

Finally, we compare the results obtained from human traders in these three markets against computer simulations of markets populated with minimal intelligence (MI) and adaptive learning (AL) algorithmic agents described in Section 4. These computer simulations provide dynamic bases of comparison for markets populated by profit motivated human traders. We simulate each of the three market structures 1000 times with specified algorithmic traders, and present the results in quintile bands of gray (wherever appropriate) to serve as the background for easy visual comparison with the theoretical equilibrium benchmarks outcomes of markets populated by profit-motivated human traders in the foreground in the figures.

<sup>12</sup> The results for the non-cooperative equilibrium are delicately dependent on the formulation of details of the game; see Shubik (1959), Wilson (1978), and Schmeidler (1980). In some models it is possible that there is no pure strategy non-cooperative equilibrium, in others there may be a multiplicity of equilibria with the same value.

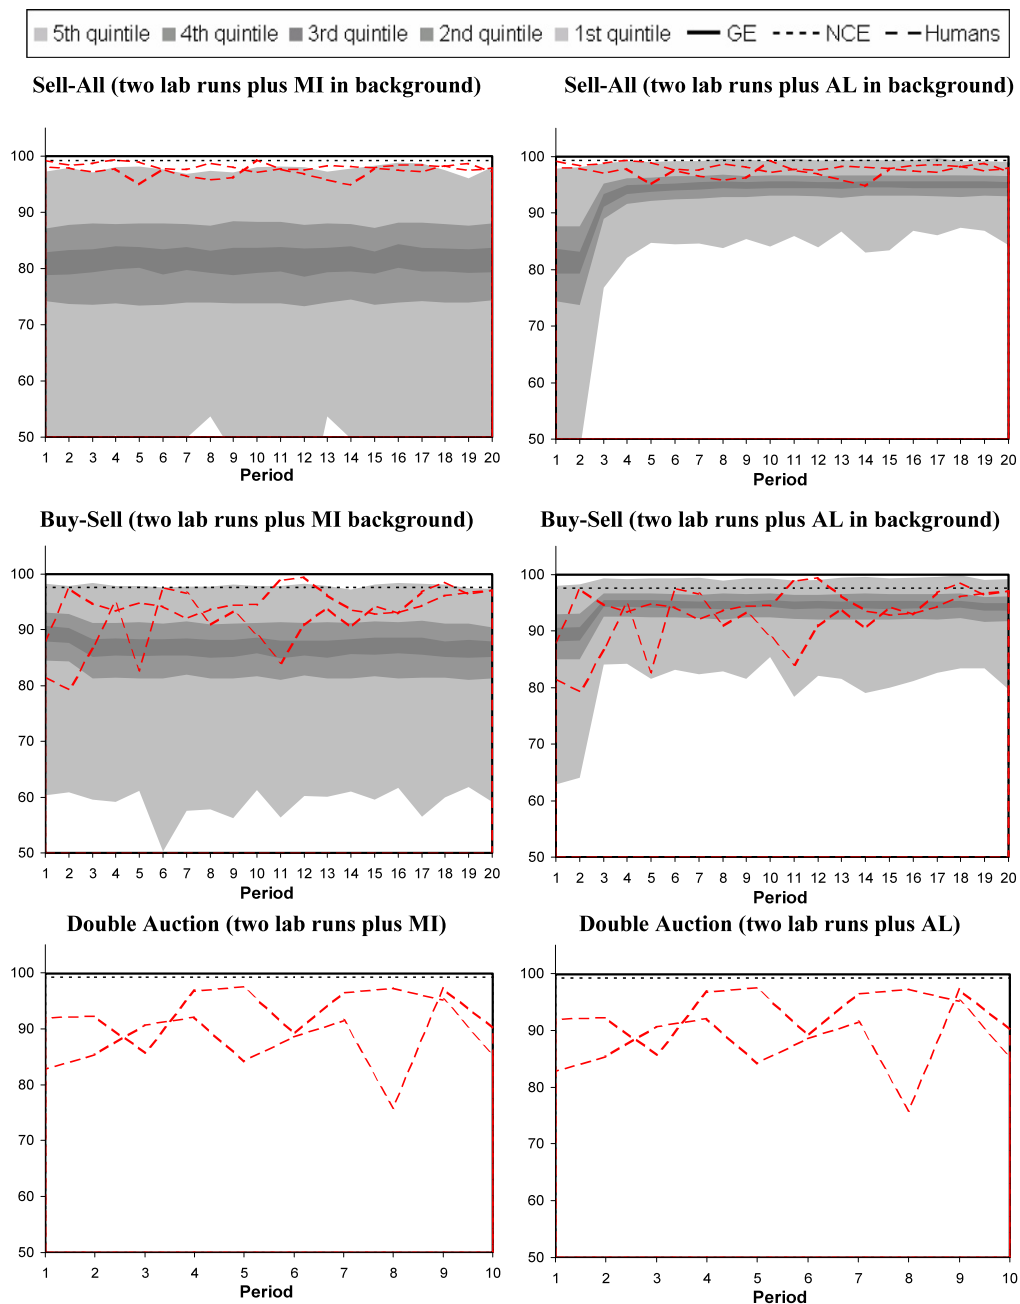


Fig. 1. Efficiency of allocations (average earnings).

Each market statistic observed over the 1000 replications is sorted into quintiles for each period. Bands in shades of gray in the background of Figs. 1 to 8 show the distribution of the performance of the markets under the specified trading algorithms. Note that in Figs. 1, 3 and 6, the double-auction simulations have zero dispersion and the quintiles collapse the bands of gray to zero width.

The discussion of results (shown in Tables 5 and 6, and Figs. 1–8) is organized around five conjectures.

*6.1. The non-cooperative and general competitive equilibrium models provide a reasonable anchor to locate most (but not all) the observed outcomes of the three market mechanisms*

An assessment of how well the CGE and NCE models organize the empirical observations of the six above mentioned measures of market performance can be seen in Figs. 1 to 7. Efficiency of the three markets (Fig. 1) approaches the predic-

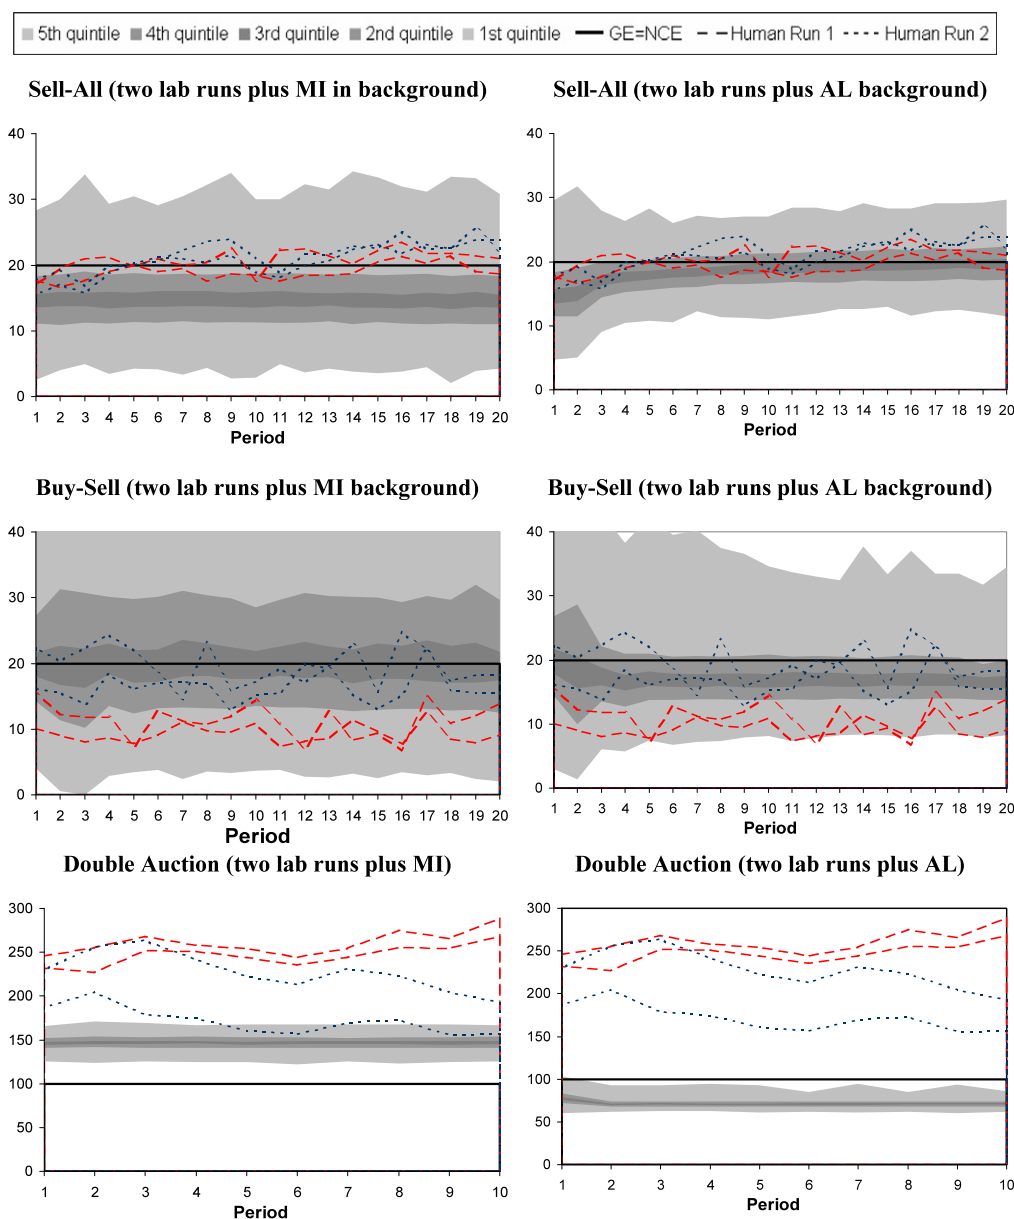


Fig. 2. Average transaction prices of goods A and B in the lab and in the MI and AL simulations (quintiles of distribution of 1000 runs).

tions of CGE (100 percent) and NCE (99.5 percent, 97.6 percent, and 99.5 percent for sell-all, buy-sell, and double auction respectively). Prices (Fig. 2) in sell-all markets are clustered around the joint CGE-NCE prediction of 20, but deviate significantly in buy-sell and DA. Symmetry of allocations to owned and non-owned goods (Fig. 3) is clustered around the NCE prediction (0.64) in buy-sell markets, but remains below the NCE prediction of 0.82 in sell-all and DA. Symmetry falls significantly short of CGE prediction of 1.0 in all markets. Unspent money (Fig. 4) is clustered around the joint CGE-NCE prediction of 32.92 percent in sell-all, and remains at or above the NCE prediction of 60.98 percent in buy-sell (this measure of performance is undefined for DA). Cross-sectional standard deviation of earnings as percentage of CGE earnings (100%, see Fig. 5) is less than 25 percent in sell-all, and in 25–50 percent range in buy-sell and DA (as compared to the zero prediction of the two models, both being non-stochastic).

Goods traded as a percentage of the volume needed to achieve CGE (Fig. 6) are in the proximity of the CGE and NCE predictions. Fig. 7 shows that these results remain qualitatively unchanged when the number of traders is changed from 5 + 5 to 10 + 10 in buy-sell markets.

On the whole, we conclude that the two equilibrium models are reasonable but far from perfect candidates to serve as the domains of attraction of the three market institutions examined here.

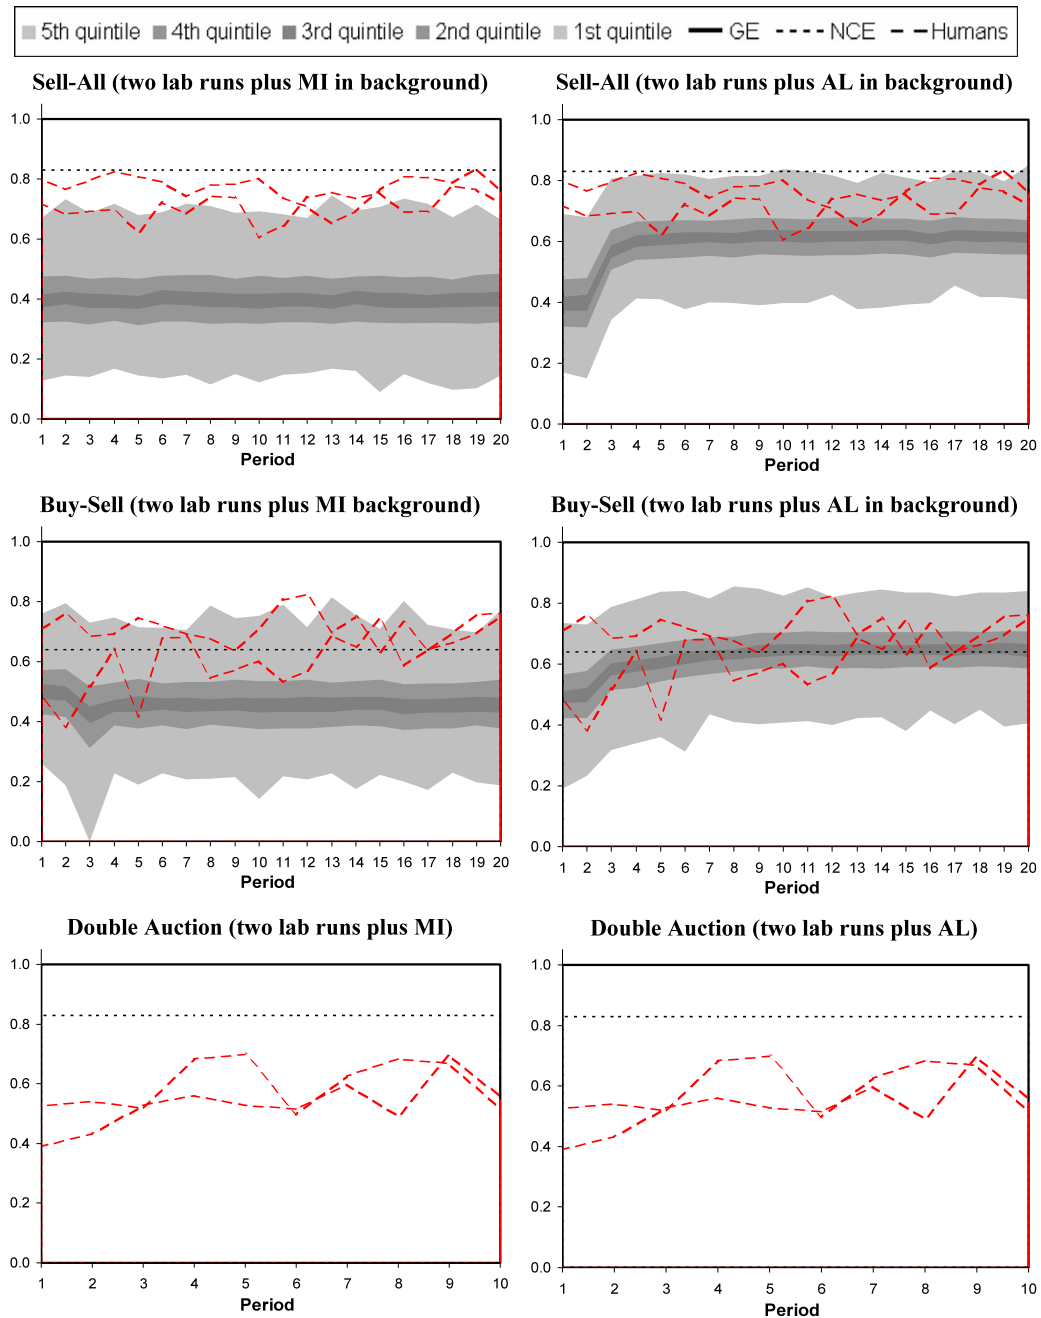


Fig. 3. Symmetry of allocations.

## 6.2. There is some evidence that outcomes tend to get closer to CGE predictions as the number of players increases

We repeated the 5 + 5 subject buy-sell market with 10 + 10 subjects to examine the direction of effects of increasing the number of traders. This effect can be assessed by comparing the six panels of Fig. 7 with the corresponding buy-sell panels of Figs. 1 to 6. In Fig. 7, efficiency is closer to 100 percent, prices are less volatile and closer to 20 on average (16.42 with 20 traders vs. 14.16 with ten traders), symmetry is higher and above 0.8, unspent money is slightly closer to 50 percent and less volatile, standard deviation of earnings declines over 20 periods to less than 10 percent, and the goods traded as a percentage of what is needed to reach CGE is also less volatile and closer to 100 percent. Since, we did not conduct sell-all and DA markets with 10 + 10 traders, this evidence of closer convergence to CGE with the increase in the number of agents in buy-sell markets is strongly suggestive, but not conclusive.

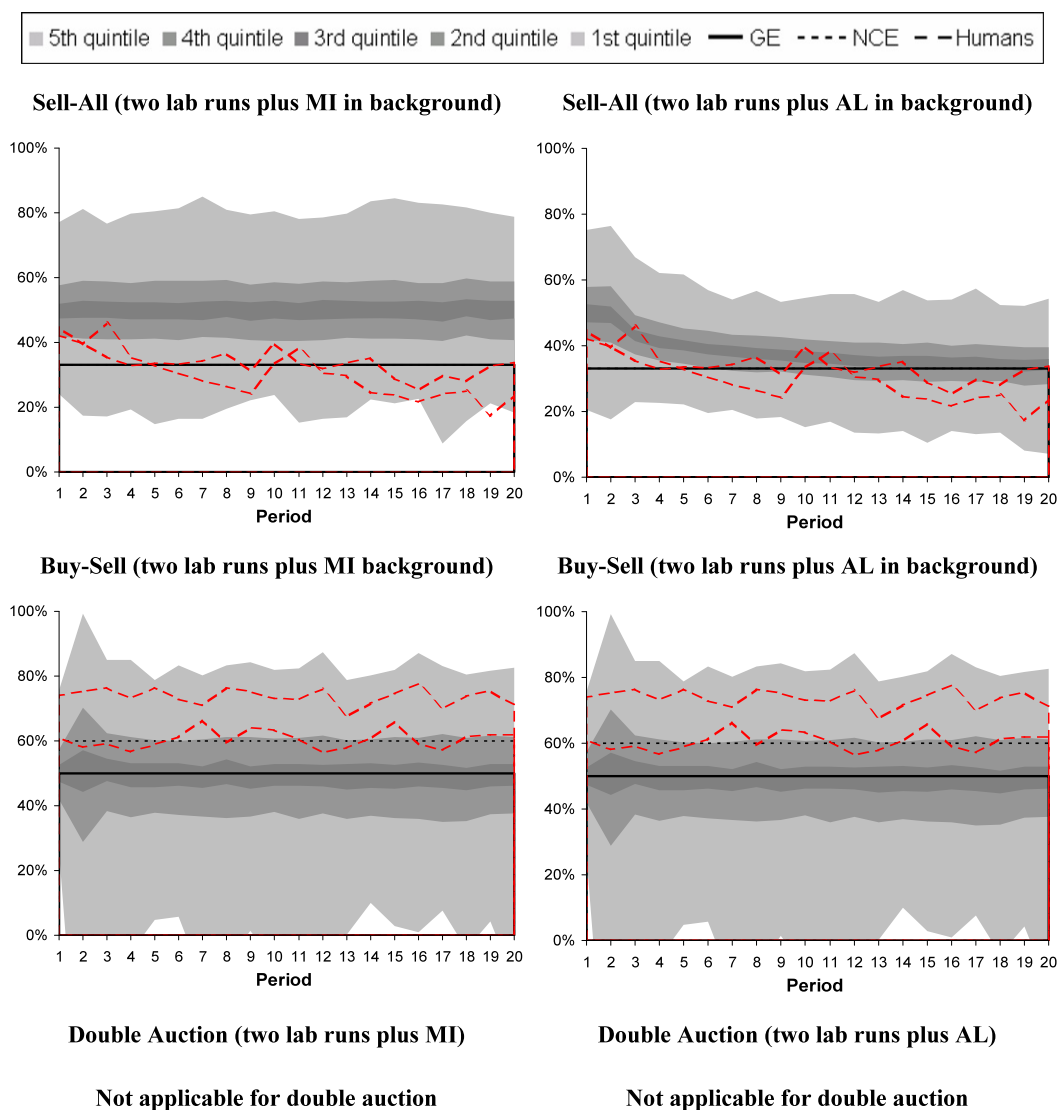


Fig. 4. Unspent money as percentage of initial endowment.

Table 4

Equilibrium predictions for the three market games.

Runs	Market game	Autarky	General equilibrium	Non-cooperative equilibrium
1 + 2	Sell-all	$P_A = P_B = NA$ $X_A = 200$ or 0 $X_B = 200$ or 0 Net money = 0 Points = 0	$P_A = P_B = 20$ $X_A = X_B = 100$ Net money = 0 Points = 1000	$P_A = P_B = 20.12$ $X_{own} = 110$ $X_{other} = 90$ Net money = 0 Points = 995
3 + 4	Buy-sell	$P_A = P_B = NA$ $X_A = 200$ or 0 $X_B = 200$ or 0 Net money = 0 Points = 0	$P_A = P_B = 20$ $X_A = X_B = 100$ Net money = 0 Points = 1000	$P_A = P_B = 20$ $X_{own} = 122$ $X_{other} = 78$ Net money = 0 Points = 976
5 + 6	Double auction	$P_A = P_B = NA$ $X_A = 20$ or 0 $X_B = 20$ or 0 Net money = 0 Points = 0	$P_A = P_B = 100$ $X_A = 20$ or 0 $X_B = 20$ or 0 Net money = 0 Points = 1000	$P_A = P_B = 100$ $X_A = 11$ $X_B = 9$ Net money = 0 Points = 995

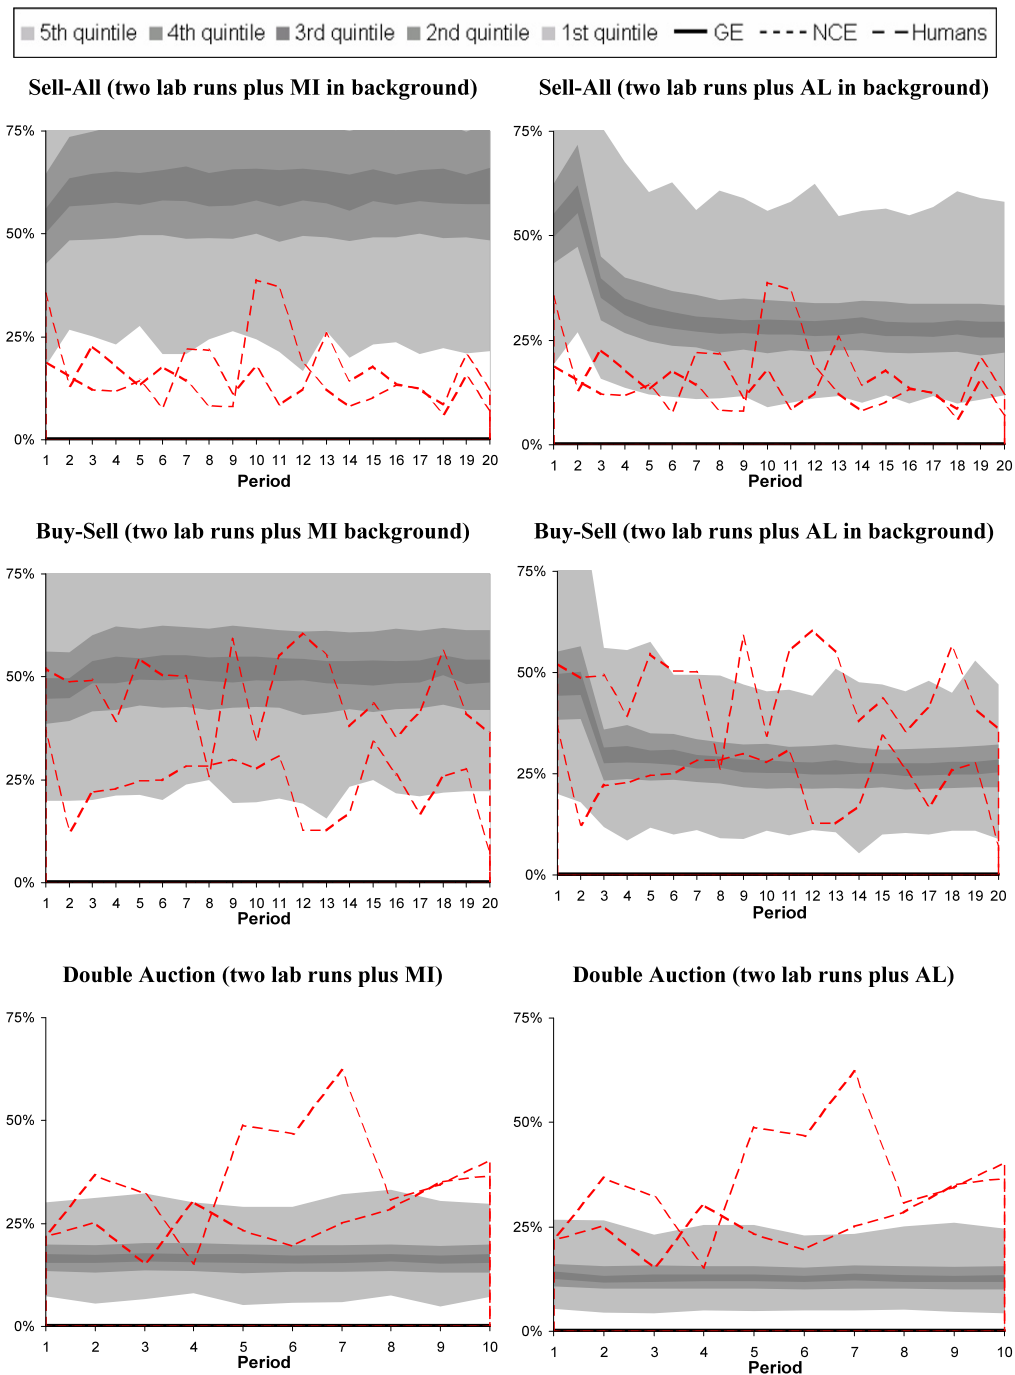


Fig. 5. Standard deviation of individual earnings per period.

### 6.3. The outcome paths from the three market mechanisms exhibit significant and persistent differences among them

A comparison across three rows of panels in Figs. 1, 2, 3, 5, 6, and 7 (and across two rows in Fig. 4) reveal significant and persistent differences in the paths of outcomes in these three kinds of markets. Efficiency ranks highest and most stable in sell-all markets, followed by buy-sell and DA (Fig. 1). While the volatility of prices across the three markets is comparable, the location of prices in sell-all is at, in buy-sell is above, and in DA is below, the CGE predictions. Symmetry of allocations to owned and non-owned goods is closest to 1.0 in sell-all, followed by buy-sell and DA, respectively, at lower levels. Unspent money in buy-sell is systematically higher than in sell-all (it is undefined in DA). The cross sectional

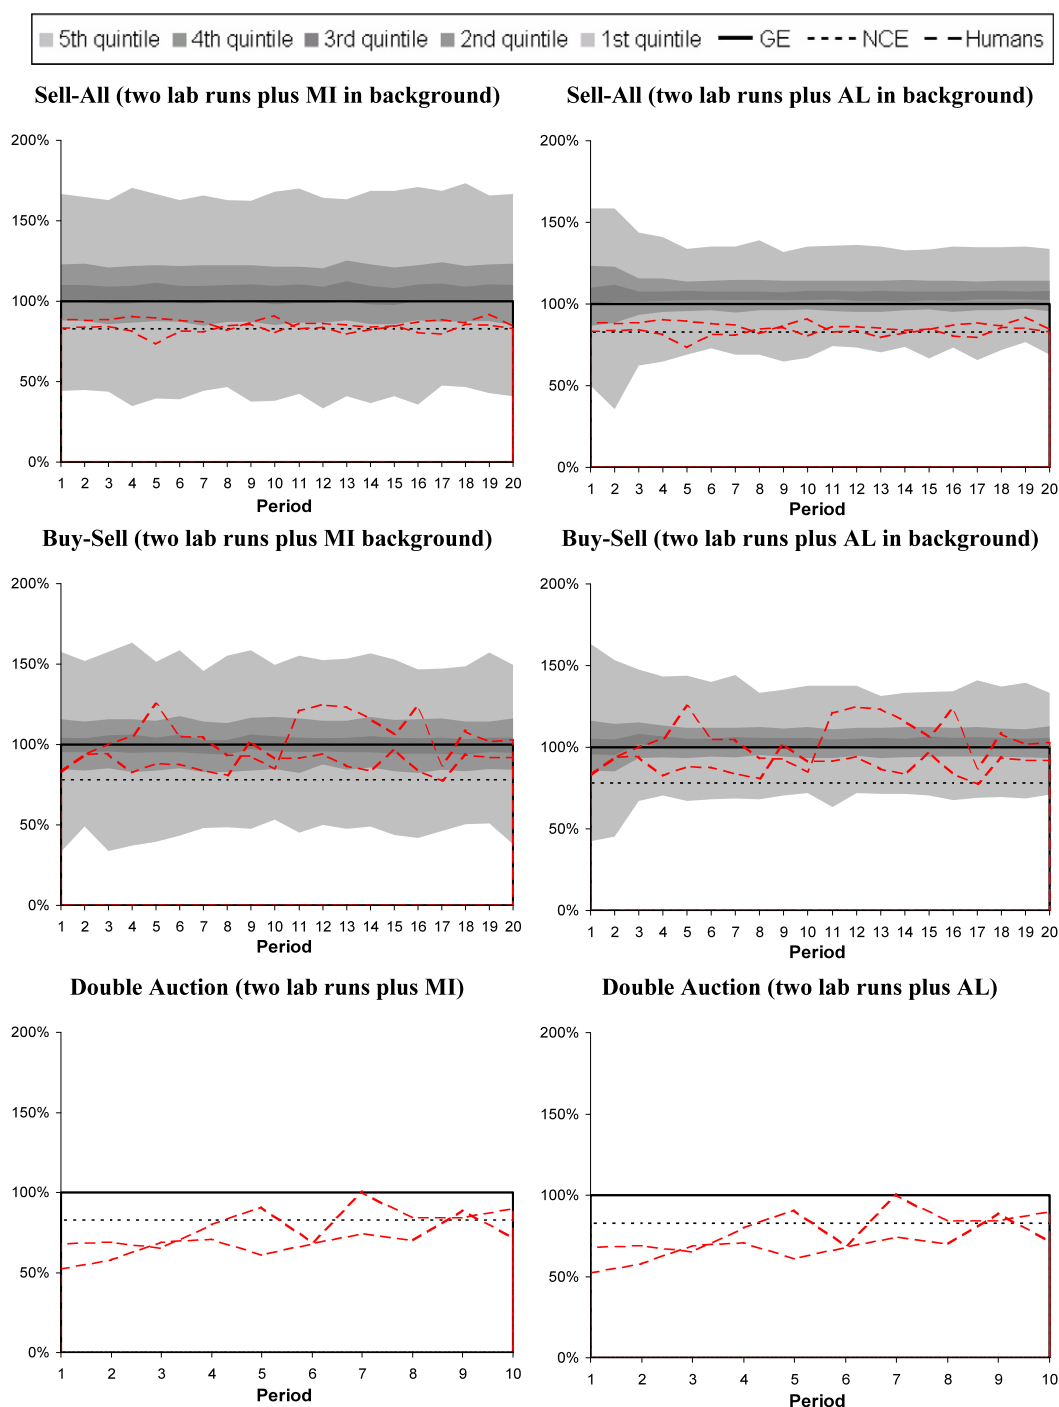
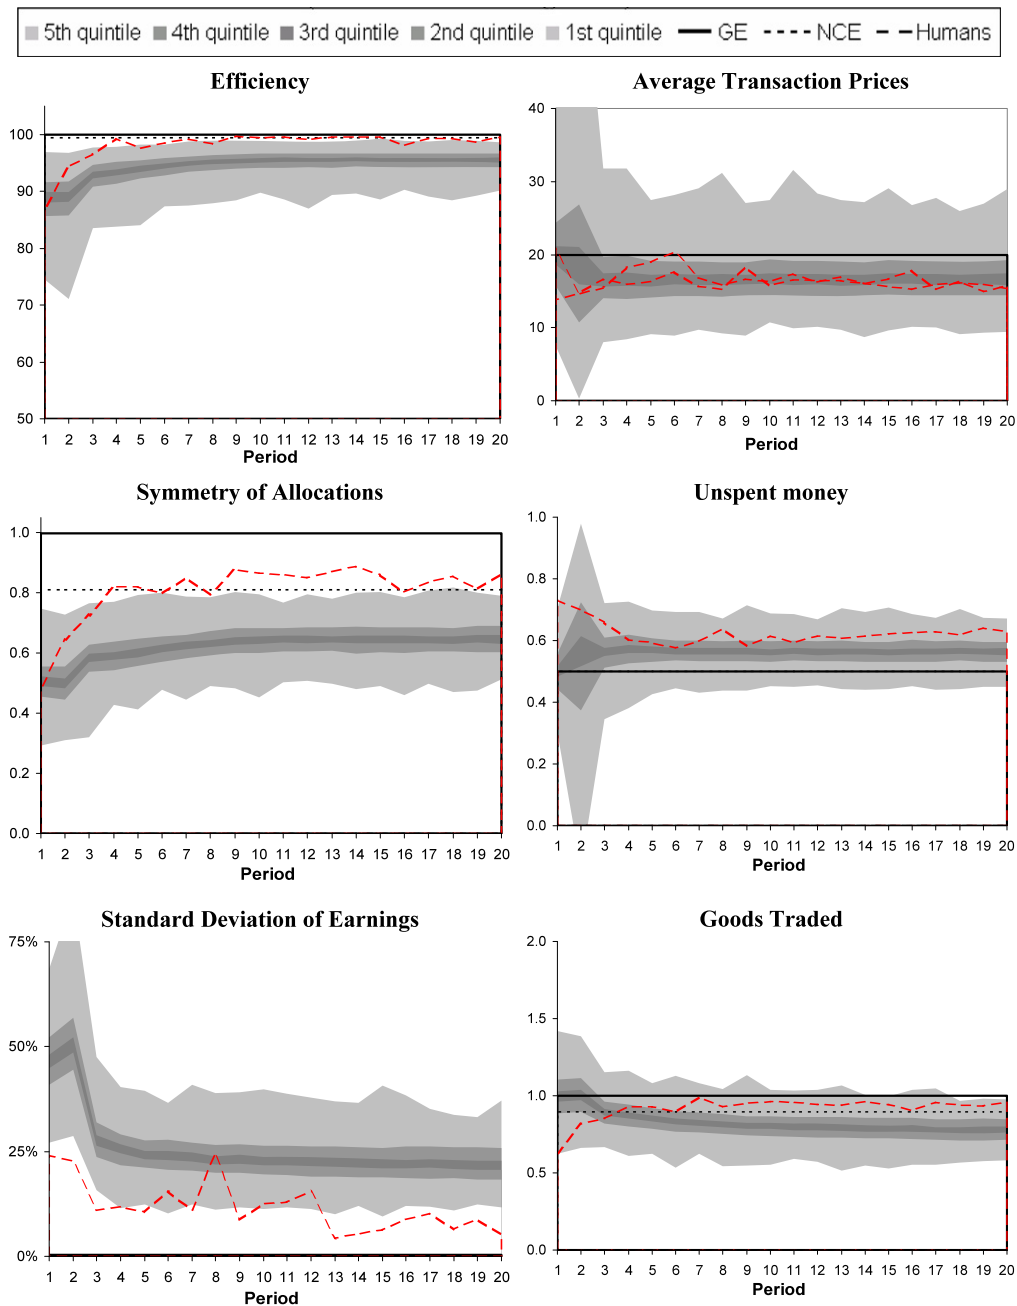


Fig. 6. Goods traded as percentage of trade needed to achieve CGE.

standard deviation of individual earnings in sell-all is about one half of what is observed in buy-sell and DA. Finally the volume of goods traded in sell-all and DA is systematically lower than in buy-sell markets.

These results support the prior experimental findings that the rules by which agents are allowed to participate, and price and allocations determined, affect the outcomes of markets. Although the CGE and NCE models do point to the general domain of attraction for outcomes of all three markets, the experiment reveals significant and persistent differences among the outcome paths observed in the three market mechanisms. In abstracting away from the variations in market rules, one must also forego the opportunity to identify and understand the systematic consequences of these variations.




**Fig. 7.** Overview data for buy-sell market with  $n = 20$  traders (1 lab run + AL background).

**6.4. Unlike the well-known results from many partial equilibrium double auctions, prices and allocations in double auctions with full feedback reveal significant and persistent deviations from CGE predictions**

Beginning with the well-known results reported by Smith (1962), the outcomes of double auctions in partial equilibrium settings have been found to lie remarkably close to the predictions of theory in most circumstances. Our experiment, conducted with full feedback of general equilibrium conditions shows the market outcomes of DA to be less reassuring. As discussed above, compared to sell-all and buy-sell, the performance of DA markets is no closer to (and often further away from) the respective CGE predictions in any of the five dimensions (efficiency, prices, symmetry of allocations, dispersion of individual earnings, and trading volume). Moreover, in absolute terms, the outcomes of DA in this general equilibrium setting is worse than the outcomes generally observed in laboratory experiments in partial equilibrium settings of which Smith (1962) is a good example. We conjecture that the fact that we had two DA-markets (compared to one often used in





**Fig. 8.** Double auction transaction price paths within individual trading periods with MI and AL traders (grey lines show individual runs, the dark line with diamonds the average).

the literature) and that subject need to conduct 20 transactions to reach CGE (compared to only two numbers they need to get right in sell-all and buy-sell) may explain why DA does relatively poorly in our experiments.

*6.5. The dynamics of markets populated by profit-motivated human subjects is at least partially captured in markets populated by simple algorithmic traders, supporting the importance of market structures in determining their outcomes*

This can be seen in the comparison of the performance of these markets populated by profit motivated human traders (marked by dashed lines in the figures) to their performance when they are populated by two different kinds of simple algorithmic traders (minimally intelligent or MI and adaptive learning of AL, see Section 4 and Appendix B for descriptions) depicted in quintiles of 1000 replications of simulations in shades of gray in the same figure. Given the strong assumptions of rationality used to derive theoretical equilibria, one might have expected that the performance of markets populated by such simple agents would fall far short of the performance of markets populated by human subjects. In many respects it does. Yet, it is worth examining how well the market institutions perform even when they are populated by such simple agents—minimally intelligent agents do not optimize, have no memory and no learning, while the adaptive learning agents have simple memory and rudimentary learning.

**Table 5**

Market data on the two double auction markets.

		Goods in market	Money in market	Goods traded	Money paid	Turnover stocks	Turnover money	Transactions/trader/period
Human	Run 5	200	40,000	994	252,362	5.0	6.3	19.9
	Run 6	200	40,000	1114	214,716	5.6	5.4	20.3
Simulations	MI	200	40,000	1092	160,747	5.5*	4.0*	21.8
	AL	200	40,000	1216	87,309	6.1*	2.2*	24.3

\* Median over 1000 replications of the market.

Median efficiency (see Fig. 1) of sell-all markets with MI agents is lower than with human agents but still around 80 percent (compared to upper 90s for humans). When MI agents are replaced by AL agents with a little learning, median efficiency of sell-all markets jumps up sharply within the first four periods and stabilizes in mid-90s. In buy–sell markets, efficiency with humans is in low 90s and the median efficiency with MI traders is in high 80s. Replacement of MI by AL results in a sharp jump within the first few periods to mid-90s—a level equal to or higher than the level achieved by human subjects. In DA panels at the bottom of Fig. 1, there are no gray quintile bands because the this market always achieves the maximum efficiency when populated with algorithmic traders dominating the efficiency of this market with human traders; in the amount of time they were allowed to trade, the algorithmic traders extracted all possible surplus.

With respect to transaction prices (see Fig. 2), the performance of the market institutions populated by algorithmic traders is as good, if not better, than with respect to efficiency. In sell-all markets, median transaction price of about 15 observed with MI agents is significantly lower than the CGE prediction 20 around which human agent prices are located. However, as with efficiency, replacement of MI by AL agents quickly raises the median transaction prices close to the CGE prediction of 20. In buy–sell markets, median transaction prices are equal to CGE with MI, and learning of AL agents moves these prices away from CGE, but farther out than what is observed with human agents. In double auctions, median prices with MI and AL agents lie much closer to CGE predictions than with human agents. In fact (as seen in Fig. 8) the transaction prices in DA converge asymptotically to CGE, and the median prices deviate from CGE only because with algorithmic traders, this convergence occurs gradually over many transactions within each period.

With respect to symmetry of allocations (see Fig. 3), MI agents achieve lower results (about 0.4) in sell-all and buy–sell than human agents, and the AL agents rapidly improve the symmetry to around 0.6 within the first few periods. In DA markets, both MI as well as AL agents easily beat human markets by achieving a perfect 1.

With respect to unspent money (see Fig. 4), MI leave more than CGE (and human agents) money unspent (50 as compared to 33 percent) in sell-all markets. Replacement of MI by AL agents progressively eliminates this excess saving and the median converges asymptotically to the CGE level. In buy–sell markets, the median money unspent is close to the CGE prediction of 50 percent and much lower than human agents. Here, replacement of MI by AL agents and the consequent learning raises the median unspent money to a level above the CGE prediction, but it is still lower than what is achieved by human agents. The reason for the increase is that AL agents look at their past highest earnings and “learn” from them—high earnings usually occur when they sell few of their endowed goods and do not overspend on the other good—thus the amount unspent is above 50 percent.

The median of cross-sectional standard deviation of individual earnings (as percent of earnings in CGE = 1000 points) in sell-all markets with MI agents is high in mid-sixties percent compared to 15–20 percent with humans. Replacement of MI by AL agents lowers the median standard deviation to about 25 percent which is still higher than in markets populated with human traders. Buy–sell markets also show a lower standard deviation with human agents than with MI, but the median performance of AL agents is equal to or better than humans. In DA, the median performance of both MI as well as AL agent markets is better than the human markets.

Finally, with respect to the trading volume observed as a percentage of trading volume needed to reach CGE, the median performance of MI and AL markets is close to the CGE prediction, and is as good as or better than what is observed in markets with human agents. In DA, algorithmic agents record a perfect score.

This comparative review of the performance of the three market institutions when they are populated by human and by two different kinds of algorithmic agents suggests that important aspects of Gode and Sunder (1993) results (about the significant properties of markets in partial equilibrium settings arising from their structural features as opposed to the behavior of agents who populate them) may hold in general equilibrium settings.

## 6.6. Ranking of the three market mechanisms

Competitive and non-cooperative equilibria are defined for the abstract end of the institutional spectrum of price formation processes. As we discussed in the introduction, the three minimal market institutions examined here can be located ordinally right next to this abstract end of the spectrum of market institutions. Table 6 presents the ordinal rankings of the three mechanisms with respect to their distance from the abstract end along six dimensions on the basis of how they perform when they are populated by human, and MI and AL algorithmic traders.

**Table 6**

Ranking of three market mechanisms on the basis of distance from CGE and NCE benchmarks.

Mkt. mechanism	Bench-mark	Alloc. eff.	Money balances	Symmetry	Prices	Dispersion of profits	Trading volume	Ave. rank	Ave. rank for mech.
Markets populated with human traders									
Sell-all	CGE	1	1	1	1	1	2	1.17	1.17
	NCE	1	1	2	1	1	1	1.17	
Buy-sell	CGE	2	2	2	2	3	1	2.00	2.00
	NCE	2	2	1	2	3	2	2.00	
DA	CGE	3	NA	3	3	2	3	2.80	2.80
	NCE	3	NA	3	3	2	3	2.80	
Markets populated with minimal intelligence algorithmic traders									
Sell-all	CGE	3	2	3	2	3	3	2.67	2.67
	NCE	3	2	3	2	3	3	2.67	
Buy-sell	CGE	2	1	2	1	2	2	1.67	1.67
	NCE	2	1	2	1	2	2	1.67	
DA	CGE	1	NA	1	3	1	1	1.40	1.40
	NCE	1	NA	1	3	1	1	1.40	
Markets populated with adaptive learning algorithmic traders									
Sell-all	CGE	3	1	3	2	2.5	3	2.40	2.37
	NCE	3	1	2.5	2	2.5	3	2.33	
Buy-sell	CGE	2	2	2	3	2.5	2	2.25	2.17
	NCE	2	2	1	3	2.5	2	2.10	
DA	CGE	1	NA	1	1	1	1	1.00	1.15
	NCE	1	NA	2.5	1	1	1	1.30	

While there are some deviations in six specific measures of performance, it is clear that, on the whole, when these mechanisms are populated by profit-motivated human traders, the outcomes of the sell-all mechanism is the closest to the CGE as well as NCE predictions, followed by buy-sell and double auction in that order. This ordinal ranking of correspondence to the predictions of the abstract models matches the ranking of specificity (i.e., additional assumptions) needed to define each market mechanism and also the complexity of the task required from agents. Perhaps it is not surprising that the increasing specificity of market mechanisms adds some distance between their performance and the abstract benchmark. To what degree this process will continue with additional specificity remains to be explored.

When profit-motivated human traders are replaced by MI and AL algorithmic traders, the quintile bands from 1000 replications of simulated markets suggest that the rankings of their outcomes change: DA is closest to CGE predictions (followed by buy-sell and sell-all) and buy-sell is closest to NCE (followed by sell-all and DA). This is mostly driven by the ending condition we defined for the DA: trading stopped only when any possible surplus was realized. This way efficiency, symmetry, dispersion of profits and trading volume were by definition “perfect.”

Still, we find that none of the three mechanisms dominates the other two in its proximity to the predications of the static models. In spite of the considerable differences in the cognitive capacities of human and algorithmic traders, the latter hold their own, reinforcing the Gode and Sunder (1993) results about the importance of the structural features of markets for their performance.

## 7. Concluding remarks

We report the performance of three minimal market mechanisms which are closed, full feedback models with explicit price-formation mechanisms and trade involving some form of money. The experiment reveals that (1) the non-cooperative and general competitive equilibrium models provide a reasonable anchor to locate most but not all the observed outcomes of the three market mechanisms; (2) there is some evidence that outcomes tend to get closer to CGE predictions as the number of players increases; (3) unlike well-known results from many partial equilibrium double auctions, prices and allocations in our double auctions with full feedback reveal significant and apparently persistent deviations from CGE predictions; (4) the outcome paths from the three market mechanisms exhibit significant and persistent differences among them; (5) since the dynamics of markets populated by profit-motivated human subjects is at least partially captured in markets when humans subjects are replaced by simple algorithmic traders, the importance of market structures in determining their outcomes is reinforced; and (6) none of the three markets necessarily dominates the others across the six measures of performance.

The comparison of market mechanisms populated by profit-motivated human traders with the frequency distribution (quintiles) of those populated by MI (minimally intelligent) and AL (adaptive learning) agents is a methodological innovation of this paper. In contrast, the earlier work in experimental gaming focuses on comparison with static predications of various equilibrium models alone. Presenting experimental results jointly with the frequency distribution of simulations allows—in our opinion—a better understanding of both the experimental results and the simulations.

The study of these three minimal mechanisms raises a basic issue about the level of specificity/generality at which one should identify the properties of (market) mechanisms. For example, on one hand, the double auction is an obvious—and extreme—abstraction from the complex rules and design of, say, the New York Stock Exchange. If each article in its rulebook

and each feature of its design of a market helps define and determine its properties, every detail matters, and nothing can be abstracted away in the study of market mechanisms. Considered in their full details, no two markets are alike, and the study of market mechanism would constitute a voluminous encyclopedia with little generality and therefore little scientific content. On the other hand, the competitive general equilibrium models of markets abstract away the details of trading mechanisms until they are reduced to become identical. The power, and the limitations, of these models arise from their generality.

This paper takes only a small step away from total generality by considering three minimally specific trading mechanisms—sell all, buy–sell, and double auction—in forms which are still highly abstract relative to what we see in the world of trade and commerce. We find that introducing even a small amount of specificity differentiates the paths markets take towards the general equilibrium prediction. It seems reasonable to conjecture that additional specificity in market mechanisms may reveal further differentiation in their properties, albeit at a diminishing rate and importance.

If the properties of mechanisms depend on the level of specificity/generality at which we study them, what is the appropriate level for their use? This question is not unique to economics and is shared with other sciences. Boyle's Law (pressure  $\times$  absolute temperature = a constant) for gases, and Ohm's Law (voltage/current = a constant) for electricity are so powerful and simple in their generality, and yet must be modified to specific gases and circuits in most practical applications. A science consists of a spectrum of laws that extend from most general approximations at one end to increasingly specific details at the other where it blends into engineering. The appropriate level of detail and specificity can be determined only from the question sought to be answered through the investigation.

As social institutions, mass market mechanisms may have evolved to minimize the importance of individual social psychological factors and the experiments presented here support this observation. They also suggest that the non-cooperative equilibrium approach is more fundamental than the competitive equilibrium, with the former encompassing the latter as a special limiting case. Furthermore the former requires the full specification of price formation mechanisms and the simplest of such mechanisms are studied here.

An important question, both in theory and in experimentation has been raised here in the treatment of terminal value of money to the experimental subjects. Theory requires that terminal or "salvage value" conditions be imposed if the game has a finite termination. Furthermore in many formal economic models a discount factor plays an important role. Yet our runs indicated that for the most part human players pay little attention to terminal conditions until close to the very end. In further experimentation it appears to be highly desirable to devise an appropriate control to study this phenomenon.

## Appendix A

### A.1. Notation

$b_j^i$  = the bid of individual  $i$  ( $i = 1, \dots, n$ ) in market  $j$  ( $j = 1, 2$ )  
 $\alpha$  = utility function scaling parameter, the same for each trader  
 $p_j$  = price of commodity  $j$   
 $M$  = initial money holding of each trader  
 $(a, 0)$  = initial holding of goods of type 1  
 $(0, a)$  = initial holdings of goods of type 2

### A.2. Calculations for sell-all

An individual  $i$  initially endowed with good  $j$  wishes to maximize his payoff function which is of the form:

$$\pi^i = \alpha \sqrt{\frac{b_1^i b_2^i}{p_1 p_2}} + (M - b_1^i - b_2^i + p_j a).$$

The calculation for the sell-all model requires to solution of the two equations derived for each trader from the first order conditions on the bidding in the two goods markets. By symmetry we need only be concerned with one type of trader.

We obtain the equation

$$\frac{b_2}{b_1} \left( \frac{(n-1)b_1 + nb_2}{nb_1 + (n-1)b_2} \right) = \frac{n}{n-1}.$$

As  $n$  becomes large this yields  $b_1 = b_2$ . Substituting in for  $b_1$  in terms of  $b_2$  we can calculate Table 1.

### A.3. Calculations for buy–sell

The payoff function for Player 1 in the buy–sell market is given by

$$\pi = \alpha \sqrt{\frac{b_1^1}{p_2}} (a - q_2^1) + (M - b_1^1 + p_2 q_2^1).$$

And similarly for Player 2; where  $q_j^i$  is the amount of good  $j$  offered for sale by individual  $i$  in market  $j$ .

We obtain from individual maximization of these equations the following values

$$b = \frac{\alpha a(n-1)^2}{2(n^2 + n - 1)}, \quad q = \frac{a(n-1)^2}{2(n^2 + n - 1)}.$$

These are utilized to calculate Table 2.

## Appendix B. Algorithm used for double auction with minimally intelligent (MI) and adaptive learning (AL) traders

1. Total number of traders =  $n$   
 Endowment:  $E_A/E_B/M$   
 Current balances at any point of time during trading:  $c_A/c_B/\mu$   
 Adaptive learning parameter:  $\lambda = 0.5$  (set  $\lambda = 0$  for no learning, i.e., MI algorithm)  
 Set initial price aspiration = total money endowment/total goods endowment
2. Randomly pick one of the  $n$  traders in the market with equal probability (with replacement); for the chosen trader, randomly pick one of the two markets with equal probability (with replacement).
3. For the chosen market, randomly pick bid or ask with equal probability (with replacement)
  - 3a. If bid is picked for the chosen trader for the chosen market A:  
 Calculate  $d = (100/2)((c_A + 1)c_B)^{0.5} - (c_A c_B)^{0.5}$ . Pick a uniform random number  $U \sim (\text{current bid}, \min(d, \text{price aspiration}))$ , and submit it as a bid for A.
  - 3b. If bid is picked for the chosen trader for the chosen market B:  
 Calculate  $d = (100/2)((c_B + 1)c_A)^{0.5} - (c_A c_B)^{0.5}$ . Pick a uniform random number  $U \sim (\text{current bid}, \min(d, \text{price aspiration}))$ , and submit it as a bid for B.
  - 3c. If ask is picked for the chosen trader for the chosen market A:  
 Calculate  $e = (100/2)((-c_A - 1)c_B)^{0.5} + (c_A c_B)^{0.5}$ . Pick a uniform random number  $U \sim (\max(e, \text{price aspiration}), \text{current ask})$ , and submit it as an ask for A.
  - 3d. If ask is picked for the chosen trader for the chosen market B:  
 Calculate  $e = (100/2)((-c_B - 1)c_A)^{0.5} + (c_A c_B)^{0.5}$ . Pick a uniform random number  $U \sim (\max(e, \text{price aspiration}), \text{current ask})$ , and submit it as an ask for B.
4. If the new bid is higher than the current bid, it becomes the current bid; if the new ask is lower than the current ask, it becomes the current ask.
5. Whenever current bid and current ask cross, record a transaction at price equal to current bid or current ask (whichever was submitted earlier). Adaptively adjust new price aspiration = existing price aspiration +  $\lambda^*$  (transaction price – existing price aspiration).
6. Let the simulation run for 25,000 iterations to complete a period. At the end of the period, use the final  $c_A$ ,  $c_B$ , and  $m$  for calculating earnings of each trader.
7. Repeat over the specified number of periods to complete the market.
8. Repeat over the specified number of replications of the market.

## Supplementary material

The online version of this article contains additional supplementary material.

Please visit [doi:10.1016/j.geb.2010.01.010](https://doi.org/10.1016/j.geb.2010.01.010).

## References

- Angerer, M., Huber, J., Shubik, M., Sunder, S., 2009. An economy with personal currency: Theory and evidence. Cowles Foundation Discussion paper No. 1622.
- Bertrand, J., 1883. Review of Walras' *Theorie mathématique de la richesse sociale*. *J. des Savants* (Paris) 67, 499–508.
- Cournot, A.A., 1897. *Researches into the Mathematical Principles of the Theory of Wealth*. MacMillan, New York (Bacon translation of original 1838, French).
- Dubey, P., 1982. Price-quantity strategic market games. *Econometrica* 50 (1), 111–126.
- Dubey, P., Shubik, M., 1978. The non-cooperative equilibria of a closed trading economy with market supply and bidding strategies. *J. Econ. Theory* 17 (1), 1–20.
- Dubey, P., Shubik, M., 1980. A strategic market game with price and quantity strategies. *Z. Nationalök. 40* (1–2), 25–34.
- Edgeworth, F.Y., 1925. *Papers Relating to Political Economy*, vol. I. MacMillan, London.
- Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. *Exper. Econ.* 10 (2), 171–178.
- Gode, D.K., Sunder, S., 1993. Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality. *J. Polit. Economy* 101, 119–137.
- Gode, D.K., Sunder, S., 1997. What makes markets allocatively efficient? *Quart. J. Econ.* 112, 603–630.
- Griesmer, J., Levitan, R., Shubik, M., 1967. Toward a study of bidding processes, Part IV: Games with unknown costs. *Naval Res. Logist. Quart.* 14 (4), 415–443.
- Huber, J., Shubik, M., Sunder, S., 2008. The value of fiat money with an outside bank: An experimental game. Cowles Foundation Discussion paper No. 1675.
- Huber, J., Shubik, M., Sunder, S., 2009. Default penalty as a disciplinary and selection mechanism in presence of multiple equilibria. Cowles Foundation Discussion paper No. 1730.
- Kumar, A., Shubik, M., 2004. Variation on the theme of Scarf's counter-example. *Computational Econ.* 24 (1), 1–19.

- Lim, S.S., Prescott, E.C., Sunder, S., 1994. Stationary solution to the overlapping generations model of fiat money: Experimental evidence. *Empirical Econ.* 19, 255–277.
- Lucas, R.E., 1987. *Models of Business Cycles*. Blackwell, Oxford.
- Lucas, R.E., 1988. On the mechanics of economic development. *J. Monet. Econ.* 22, 3–42.
- Lucas, R.E., Sargent, T.J., 1981. *Rational Expectations and Econometric Practice*. University of Minnesota Press, Minneapolis.
- Marimon, R., Sunder, S., 1993. Indeterminacy of equilibria in a hyperinflationary world: Experimental evidence. *Econometrica* 61 (5), 1073–1107.
- Marimon, R., Sunder, S., 1994. Expectations and learning under alternative monetary regimes: An experimental approach. *Econ. Theory* 4, 131–162.
- Marimon, R., Sunder, S., 1995. Does a constant money growth rule help stabilize inflation? Experimental evidence. *Carnegie–Rochester Conf. Ser. Public Pol.* 43, 111–156.
- Marimon, R., Spear, S., Sunder, S., 1993. Expectationally-driven market volatility: An experimental study. *J. Econ. Theory* 61 (1), 74–103.
- Milgrom, P., Weber, R., 1982. A theory of auctions and competitive bidding. *Econometrica* 50, 1089–1122.
- Nash, J.F., 1951. Non-cooperative games. *Ann. of Math. (2)* 54, 286–295.
- Plott, C., 1982. Industrial organization theory and experimental economics. *J. Econ. Lit.* 20 (4), 1485–1527.
- Quint, T., Shubik, M., 2005. Gold, fiat and credit: An elementary discussion of commodity money, fiat money, and credit: Part 2. *ICFAI J. Monet. Econ.* 3, 6–50.
- Quint, T., Shubik, M., 2007. Monopolistic and oligopolistic banking: An elementary discussion of commodity money, fiat money and credit: Part 3. *ICFAI J. Monet. Econ.* 5, 6–65.
- Satterthwaite, M.A., Williams, S.R., 1989. The rate of convergence to efficiency in the buyer's bid double auction as the markets become large. *Rev. Econ. Stud.* 56, 477–498.
- Schmeidler, D., 1980. Walrasian analysis via strategic outcome functions. *Econometrica* 48, 1585–1593.
- Shapley, L.S., 1995. *Class Notes Mimeographed*, UCLA.
- Shapley, L.S., Shubik, M., 1977. Trade using one commodity as a means of payment. *J. Polit. Economy* 85 (5), 937–968.
- Shubik, M., 1959. *Strategy and Market Structure: Competition, Oligopoly, and the Theory of Games*. Wiley, New York.
- Shubik, M., 1973. Commodity money, oligopoly, credit and bankruptcy in a general equilibrium model. *Western Econ. J.* 11 (1), 24–38.
- Shubik, M., 1999. *The Theory of Money and Financial Institutions*, vol. 1. MIT Press, Cambridge, MA; UK.
- Smith, V.L., 1962. An experimental study of competitive market behavior. *J. Polit. Economy* 70 (2), 111–137.
- Smith, V.L., 1982. Microeconomic system as an experimental science. *Amer. Econ. Rev.* 72 (5), 923–955.
- Sorin, S., 1996. Strategic market games with exchange rates. *J. Econ. Theory* 68 (2), 431–446.
- Vickrey, W., 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* 16 (1), 8–37.
- Wilson, R., 1978. Competitive exchange. *Econometrica* 46 (3), 557–585.