Sean Crockett

Learning Competitive Equilibrium in Laboratory Exchange Economies

Abstract A laboratory market for two goods is instituted to examine the hypothesis that individuals will eventually coordinate on the induced competitive equilibrium. The mechanism for exchange strongly restricts the space of agent actions, facilitating the identification of decision rules. Evidence for learning competitive equilibrium is mixed due to strong heterogeneity in decision-making. Some subjects will forego immediately available gains when they expect the market to move in a more favorable direction, a condition necessary for coordinating on the competitive outcome. However, a majority do not, and many are content to *satisfice*, though the means to do better was reasonably transparent.

Keywords Disequilibrium · Learning · Zero Intelligence · Experimental Economics

JEL Classification C91 · C92 · D50 · D83

I gratefully acknowledge the advice and tutelage of John Duffy and Stephen Spear, and would also like to thank David Grether, Dan Houser, Michael Peress, Bryan Routledge, Shyam Sunder, and participants at the Midwest Theory meetings (Bloomington) and the ESA regional meetings (Tucson) for useful feedback. Two patient and insightful anonymous referees have helped to greatly improve this paper and my future written endeavors. This project was financially supported by the Department of Economics at Carnegie Mellon University, the International Foundation for Research in Experimental Economics at George Mason University, and the Social and Information Sciences Laboratory at Caltech, as well as by Stephen Spear. Finally, special thanks is owed to Allen Geary, Jr. for co-developing the software used in these experiments, and to the Pittsburgh Experimental Economics Laboratory for use of its facilities.

S. Crockett
Social and Information Sciences Laboratory, California Institute of Technology, Pasadena, CA
91125, USA
Tel.: 1-626-395-2245
Fax: 1-626-395-2245
E-mail: scrockett@ist.caltech.edu

1 Introduction

Beginning with Smith (1962), the experimental economics literature has confirmed the descriptive relevance of the competitive model in a wide range of partial equilibrium settings, at least after several iterations of a repeated static environment. However, in recent research by Anderson, Plott, Shimomura and Granat (2004), prices in a laboratory implementation of Scarf's example track the tâtonnement limit cycle, suggesting the Arrow-Debreu-McKenzie model may be limited by more than market size in its applicability. Anderson et al. and others who have implemented general equilibrium experiments (see, for example, Nussair, Plott and Reizman (1995), Lian and Plott (1998), Williams, Smith, Ledyard and Gjerstad (2000), and Gjerstad (2004)) have typically focused on market-level analyses of double auctions. Unfortunately, it has proven difficult to study individual decision-making in this environment, in part because the state and control spaces are enormous. Identifying behavioral types in laboratory markets could prove valuable in testing and developing descriptively relevant theories of market coordination, thus helping to determine the conditions under which the Arrow-Debreu-McKenzie model applies, how quickly markets equilibrate when it does, and what limiting outcomes we might expect, if any, when it does not.

The goal of this paper is to provide insight into individual decision-making in markets. The market institution adopted presents subjects with a series of binary decisions, facilitating behavioral identification relative to the double auction, in which time is typically continuous and agents have a large set of possible actions available at any instant. In the experiments a proposed reallocation of goods is randomly generated, and each subject is given the opportunity to vote "yes" or "no" on his component of the proposal. Individual decisions are submitted, the reallocation occurs if and only if consent is unanimous, and then a new proposal is generated. This process continues for five minutes. At the end of this five-minute stage, profits corresponding to induced preferences are recorded, and then the market reopens for a new stage, starting from the original endowment of goods. The market is opened for a total ten five-minute stages, each beginning from the same endowment allocation.

Previous theoretical results motivate why this institutional framework may capture important aspects of decision-making in markets. Zero intelligence (ZI) behavior, introduced by Gode and Sunder (1993) in a partial equilibrium setting and Gode, Spear and Sunder (2004) in a 2-good general equilibrium economy, posits that individuals are uniform randomizers over actions in their utility-improving sets. Under the institution used in the present paper, the acceptance of all utility-improving trade proposals and rejection of those that diminish utility may be readily interpreted as ZI behavior, with randomization being carried out by the institution itself rather than individuals. Thus, the prevalence of ZI behavior is readily detectable in this environment.

The random proposal generator also permits analysis of how experience impacts decision-making. It is reasonable to expect that at least some subjects will adjust their behavior over time, since the partial and general equilibrium experimental literature referenced above reports strong regularities in final allocations across repetitions. By way of theoretical guidance, Crockett, Spear and Sunder (2006) (hereafter CSS) propose a learning model that puts structure on the set of

allocations an individual *will not accept*, a set that tends to get larger over repetitions. CSS demonstrate that ZI behavior coupled with their learning rule generates convergence to a competitive equilibrium in a range of environments. Because CSS represents the only general equilibrium model of adaptive behavior applicable to this pure exchange setting, the model will be reproduced in some detail and tested in this paper.

The results of the experiment align along several stylized facts. Subjects almost always traded to the contract set in each period, but evidence for learning competitive equilibrium is mixed due to substantial individual heterogeneity. Nearly one-third of the subjects demonstrated a bias towards rejecting profitable proposals if the terms of trade were poor (here “profitable” implies that the proposed reallocation would result in an induced utility increase over the present allocation). Furthermore, the aspiration level of at least 42% of this group was fluid and appeared to be updated during the course of the experiment. Thus one-third of the subjects displayed behavior appreciably more sophisticated than zero intelligence, although only 3 of 74 total subjects behaved consistently with CSS.

However, the remaining two-thirds of the subjects appeared to be much more myopic in their decision-making, and in fact 41% of all subjects behaved as stochastic ZI robots (the majority of this group were almost deterministically ZI). If such a distribution of behavioral types holds more generally, a matter of pure speculation to be sure, it would suggest that theorists should pay serious attention to the interaction of behavioral types when addressing the issue of general equilibrium coordination. The behavioral classification provided in this paper may also assist in the analysis of individual behavior in more standard market institutions like barter, call markets, and the double auction, where identification is daunting and the econometrician might benefit from more prior behavioral structure.

2 The CSS model

A finite number of agents $i = 1, \dots, M$ trade a finite number of goods $j = 1, \dots, \ell$ in a pure exchange economy. Preferences are represented by utility functions

$$u_i : \mathbb{R}_+^\ell \rightarrow \mathbb{R},$$

which are strictly increasing, strictly quasi-concave, at least twice differentiable, and satisfy the Inada condition. An *allocation* $\mathbf{x}$ is defined as an M -tuple of commodity *portfolios* $[x_1, \dots, x_M]$, where $x_i \in \mathbb{R}_+^\ell$ for $i = 1, \dots, M$, and x_{ij} is agent i 's inventory of good j . Let $\omega_i \in \mathbb{R}_+^\ell$ be the initial portfolio possessed by agent i , and ω be the initial endowment allocation. The set of feasible allocations is denoted

$$\Lambda \equiv \left\{ \mathbf{x} \in \mathbb{R}_+^{\ell M} : \sum_{i=1}^M x_i = \sum_{i=1}^M \omega_i \right\},$$

and $(\Lambda, \mathcal{A}, \mu)$ is the Lebesgue measure space of feasible allocations.

Market activity occurs in *stages*, with each stage beginning from ω and partitioned into *periods* of exchange. Let $\mathbf{x}_t^\tau$ represent the allocation held in stage τ , period t , with $\mathbf{x}_0^\tau = \omega$ for $\tau \geq 0$. The individually rational set relative to $\mathbf{x}$, restricted to feasible allocations, is denoted $IR(\mathbf{x})$, and $v_t^\tau(\mathcal{A})$ represents the likelihood that

an allocation in some feasible region $\mathcal{A}$ will be adopted in stage τ , period t . CSS place the following two restrictions on exchange.

A1 Pareto-improvement. For $t, \tau \geq 0$, $v_i^\tau(IR(\mathbf{x}_t^\tau)) = 1$.

A2 Non-intransigence. Fix $\varepsilon > 0$, and let $C^\varepsilon(\mathbf{x}_t^\tau) \in \mathbb{R}_+^{\ell M}$ be the intersection of Λ and the cube of radius ε centered on $\mathbf{x}_t^\tau$. Let γ be any subset of $C^\varepsilon(\mathbf{x}_t^\tau) \cap IR(\mathbf{x}_t^\tau)$ such that $\mu(\gamma) > 0$. Then for all $t, \tau \geq 0$, it must be the case that $v_i^\tau(\gamma) > 0$.

A1 requires that trade be Pareto-improving. A2 loosely implies that in a neighborhood of the current allocation, any individually rational Pareto-improvement can potentially be adopted. Thus no agent is permitted to be deterministically obstinate with respect to any utility-improving opportunity. CSS demonstrate that A1-2 imply that any sequence of within-stage reallocations has a Pareto optimal fixed point, and any allocation in the contract set relative to $\mathbf{x}_t^\tau$ can be attained. With regards to A2, when the support of $C^\varepsilon(\mathbf{x}_t^\tau) \cap IR(\mathbf{x}_t^\tau)$ is uniform random behavior can be interpreted as zero intelligent; such behavior within the CSS model will be called CSS-ZI.

CSS also propose a learning rule that increasingly restricts acceptable reallocations across stages. Let $\hat{\mathbf{x}}^\tau$ be the final allocation in stage τ . Because the allocation is Pareto optimal, the normalized utility gradients for all agents will be equal. Call this normalized vector p^τ . Now, define the i^{th} agent's gain in stage τ as

$$\lambda_i^{\tau+1} \equiv p^\tau \cdot (\hat{x}_i^\tau - \omega_i),$$

and for convenience let $\lambda_i^0 \equiv -\infty$. CSS present the following learning rule:

A3 Decreasing subsidization. Suppose $T \geq 0$ is the current stage, and let $\bar{\tau}_i$ be the maximum $\tau \in [0, T]$ such that $\lambda_i^{\bar{\tau}_i} < 0$. Then $\mathbf{x}_t^T$ can be adopted in period t of stage T only if $p^{\bar{\tau}_i} \cdot (x_{i,t}^T - \omega_i) \geq \lambda_i^{\bar{\tau}_i}$ for all i .

At the end of stage τ , if the gain $\lambda_i^{\tau+1} < 0$ (that is, if the “gain” is in fact a loss), agent i is said to be *subsidizing* other agents. A3 requires each agent to provide a smaller subsidy in each successive stage (under her present constraint price $\bar{\tau}_i$). If no agent is providing a subsidy, the economy has reached a competitive equilibrium, since $\lambda_i^{\tau+1} \geq 0$ for all i implies

$$\sum_{i=1}^M p^\tau (\hat{x}_i^\tau - \omega_i) = p^\tau \cdot \sum_{i=1}^M (\hat{x}_i^\tau - \omega_i) \geq 0.$$

From the restrictions on preferences, it must be the case that $p^\tau \gg 0$, while, from feasibility, $\sum_{i=1}^M (\hat{x}_i^\tau - \omega_i) = 0$. Thus $p^\tau \cdot (\hat{x}_i^\tau - \omega_i) = 0$ for all i , and the economy has reached a competitive equilibrium. CSS proves that an exchange process satisfying A1-3 will converge to a competitive equilibrium when there are two agents, and conjecture that generic convergence might be obtained for larger economies (they also provide numerical evidence of convergence in Scarf's example and 3-agent Cobb-Douglas economies).¹

¹ When there are more than two goods it is necessary to strengthen A3 to rule out constraint cycling. Since only two goods appear in this experiment A3 will suffice.

Fig. 1 Across-Stage Learning in a 2×2 Economy

Geometric intuition is useful to motivate this learning process. Agent i 's binding subsidization constraint in A3 is $\lambda_i^{\bar{\tau}_i}$. A3 requires future allocations to lie in the halfspace above the hyperplane passing through $\hat{\mathbf{x}}^{\bar{\tau}_i}$ and orthogonal to the (common) utility gradient at that allocation. Interpreting a normalization of this gradient as a vector of prices, the constraint hyperplane can be viewed as a budget set through $\hat{\mathbf{x}}^{\bar{\tau}_i}$ at which the endowment is not affordable. Therefore, if the prices on which all agents implicitly agree at the end of stage $\bar{\tau}_i$ (by virtue of Pareto optimality) had been in effect at the beginning of the stage, agent i could have afforded to reach a higher indifference surface by waiting until the end of the stage to trade. The agent interprets subsidization as a signal that she might have secured better terms of trade (and presumably greater utility), and A3 reflects a simple heuristic designed to reduce such missed opportunities in the future.

In Figure 1, at some stage-ending allocation $\mathbf{x}$ in a 2×2 economy, agent 1 was a subsidizer (note the initial endowment resides in the half-space above her constraint hyperplane). A1-A3 imply that the final allocation in the following stage must lie in the region of the contract set above this hyperplane. Suppose in the following stage the allocation $\mathbf{y}$ is implemented. At this allocation agent 2 is a subsidizer, so subsequent allocations must lie above her constraint hyperplane. Since agent 1 is a subsidy receiver in this stage, her subsidization constraint remains unchanged from the first stage. Thus potential reallocations in the subsequent stage are restricted to the subset of the contract set contained in the shaded region of the figure. And so forth.

3 Experimental Design

Subjects were recruited through the University of Pittsburgh Experimental Economics Laboratory (P.E.E.L.) during academic year 2003-4. Most subjects were

undergraduates enrolled at the University of Pittsburgh or Carnegie Mellon University. Each was paid \$5 for completing the experiment and additional performance-based compensation. Specifically, subjects were awarded points at the end of each stage and paid the sum of these stage-ending points divided by 100, in dollars, at the end of the session. Subjects were presented with two minutes of verbal instruction to highlight payment procedures, and then familiarized themselves with the trading interface by completing an on-line tutorial, which can be downloaded from http://www.ist.caltech.edu/~scrockett/cssexp_instr_0404.ppt. Subjects were permitted to complete the tutorial at their own pace, and signalled the experiment supervisor when they were ready to begin. They also received printed instructions for convenient reference during the experiment.

Each subject participated in a two agent, two good economy. Each economy consisted of ten stages, where each stage was a fixed length of five minutes and consisted of a series of feasible proposed reallocations. A subject was presented with each proposed new portfolio and asked to accept or reject it. Trade occurred if and only if consent was unanimous. Stage length, the number of stages, and the result of group decisions were common knowledge, but individual endowments, induced preferences, and individual decisions were private information (of course, the decision of a subject's trading partner to accept or reject a proposal could be deduced if and only if said subject accepted the proposal). The aggregate endowment of each good in the economy was 1,000 units.

Allocations were restricted to an integer grid sufficiently fine to permit each of 1,000 CSS-ZI simulations to converge to an approximate competitive equilibrium. Specifically, after normalizing the distance from the initial endowment to the competitive equilibrium to be equal to one, each of these simulations reached an allocation within a distance of 0.01 of the competitive equilibrium in the tenth stage. Further, all Pareto-improvements available by making the space continuous at each simulated final allocation would have generated less than one penny of additional profit for each agent under the payoff scale adopted in the actual experiment. An average of 60-80% of the proposals received by subjects in the first three treatments occurred *after* the economy had already reached a near-Pareto optimal allocation. Thus stages were sufficiently long for most subjects to be selective with regards to which reallocation proposals they were willing to accept and still exhaust gains from trade.

The experiment was designed to answer some basic questions about individual decision-making in markets for multiple goods: Is behavior individually rational? Are individuals generally willing to take all utility-improving opportunities presented to them or will they occasionally forego some gains in hope of consummating larger ones? When the latter description applies, do individuals adjust their behavior over time according to market conditions? If so, is CSS a good description of the adaptive process, or can other processes be identified? The simplest environment in which to answer these questions is a two agent, two good economy. More goods would make the task of inducing preferences much more difficult. An increase in the number of subjects would require longer stages or a decrease in the rate of rejection by subjects in order to implement Pareto optimal outcomes, because the adopted institution grants veto power over reallocations to any individual rejector. Thus all else being equal, a two-agent, two-good economy is best suited for the exhibition of learning in this institution. That is, in this par-

ticular institution if subjects don't exhibit learning in the 2×2 case, then there is a strong reason to believe they would not exhibit learning in a more general economy.

3.1 The trade proposal generation process

Each stage consisted of a series of random proposed reallocations. Most draws were from a uniform random distribution on a 200×200 unit square centered on the current allocation, constrained to be feasible and non-gifting (that is, each subject was asked to part with units of one good to receive units of the other). However, occasionally proposals were 'directed' to be Pareto-improving. The first proposal of a stage was always uniform random, as was the first proposal following any exchange and the first proposal following a Pareto-improving proposal (regardless of whether or not it was adopted). After each subsequent failure to generate a Pareto-improving proposal or consummate a trade, the probability that the next proposal would be randomly chosen from intersection of the Pareto set with the 200×200 square was increased by 0.2. Thus, after five non-Pareto-improving proposals were consecutively rejected the sixth proposal was necessarily Pareto-improving.

The directed proposals increased the opportunity for subjects to exhaust gains from trade. Near the Pareto set, the proportion of Pareto-improving to total allocations was typically quite small; under a uniform random proposal generator, subjects might see scores of consecutive proposals that were not Pareto-improving. If directed proposals were not occasionally generated, then many stage-ending allocations would not have been sufficiently near-Pareto optimal to effectively apply the CSS learning rule,² or it would have been necessary to make stages so long that subjects would have become excessively bored. Exerting additional control over the proposal generation process was essential for behavioral identification and testing CSS.

A potential problem with using directed proposals is that the data generating process was presented ambiguously to subjects. If only the uniform random generator were used, subjects could be briefed on the process and then form objective expectations with respect to the location of subsequent proposals. Instead, subjects were simply informed they would face a series of proposed reallocations, with no additional information about how the proposals would be generated. However, it is important to note the theories to be tested in this paper (ZI, CSS) do not require such information. Further, in naturally occurring institutions like the double auction, call markets, and barter, proposals are endogenous and thus also come from a distribution generally unknown to participants who did not make them. The ambiguity introduced by directing some proposals does not itself preclude answering any of the questions that motivated this paper, whereas several questions would be difficult to answer without doing so.

A second possible criticism of directing proposals is that the market institution occasionally uses preference information that is ostensibly private. CSS itself is subject to a similar complaint because the stage termination condition utilizes

² Application of the learning rule at a non-Pareto optimal allocation can result in the competitive equilibrium being ruled out by a subsidization constraint.

private information. In the experiment this criticism is unwarranted, as the fundamental advantage of experimental vs. naturally occurring data is that the experimenter actually possesses such information and can exploit it to address questions unanswerable in the field. The purpose of the experiment is not to design a practical institution, but rather to answer questions about individual behavior which can be addressed *given* external control over preferences and the institution.

3.2 Treatments

Four treatments were developed for this experiment, each involving 16-22 human subjects. All induced preferences were a monotonic transformation of the constant elasticity of substitution (CES) function:

$$u(x_1, x_2) = \frac{\left[\left(p_1^{\frac{1}{p_3}} x^{\frac{p_3-1}{p_3}} + p_2^{\frac{1}{p_3}} y^{\frac{p_3-1}{p_3}} \right)^{\frac{p_3}{p_3-1}} \right]^{p_4}}{p_5}$$

The transformation involved exponentiating the CES function ($p_4 = 2$ in all treatments) to magnify utility differences between portfolios, and dividing the transformed function by an amount (p_5) to make the mean expected performance-based payment equal to \$10. Parameters were chosen to make the goods strong complements at the initial endowment. Complementarity generates a large individually rational set, which permitted variation in decisions of individually rational subjects. Consider the other extreme, near-linear preferences. In this case, any sequence of Pareto-improving allocations would end near the competitive equilibrium, so ZI subjects and CSS subjects would be observationally equivalent. In all treatments Player 1 began with an initial endowment of 800 units of good 1 and 200 units of good 2, while Player 2's initial endowment was 75 units of good 1 and 925 units of good 2.

In **Treatment 1** the role of Player 1 was assigned to a subject and Player 2 to a CSS-ZI computer trading program.³ Parameters were chosen so that a ZI subject would be a subsidizer most of the time ($p_1 = 0.4$ and $p_2 = p_3 = 0.6$ for the subject, and 0.75, 0.25, and 0.99, respectively, for the robot). See Figure 2,⁴ which displays the final allocations of 1,000 ZI simulations for this parametrization. Clearly, if subject i adopted a ZI trading rule in stage τ , it would likely be the case that $\lambda_i^{\tau+1} < 0$; only 5.1% of the simulations resulted in Player 1 not subsidizing. Only if a subject is more discriminating with respect to utility-improving opportunities, or becomes so over time, will clustering be observed near the competitive equilibrium. This treatment provided an opportunity to identify ZI and CSS behavior while maintaining some uniformity in the feedback received by subjects.

³ Subjects were informed they could be matched with other subjects or computerized trading programs. The trading programs were designed with a built-in decision lag of 2 to 6.5 seconds to simulate human response. This lag was slightly too long, as the average number of proposals generated in Treatment 1 was about 15% less than in Treatments 2 and 3.

⁴ In the instructions, goods were labeled X and Y rather than good 1 and 2, respectively. The figures will reference X and Y , while the text will remain consistent with previous notation and label goods 1 and 2.

Fig. 2 Final allocations in ZI simulations

Fig. 3 Median final allocations per period in CSS-ZI simulations

Treatment 2 was parameterized identically to Treatment 1, but with both player roles assigned to human subjects. This treatment was designed to facilitate the identification of non-CSS learning behavior. For example, a subject might become more selective over time without regard to his subsidization constraints, a circumstance that could be identified in the role of player 2 but not necessarily of player 1. In **Treatment 3**, $p_1 = p_2 = 0.5$ and $p_3 = 0.6$ for both subjects, which provided an opportunity to analyze how subjects behaved when each was almost equally likely to subsidize under ZI play. It should be noted that $p_5 = 7,656$ for Player 1 in this treatment, and 12,656 for Player 2. In Treatments 1 and 2, these values were 9,000 and 10,000, respectively.

The presentation of subsidization constraints was a thorny issue. It was undesirable to “coach” subjects on the issue of subsidization, but some relevant information needed to be relayed in order to test CSS. Let the marginal rate of substitution (MRS) of good 1 for good 2 at the stage-ending portfolio be called the *value* of good 1 per unit of good 2 (in the instructions it is defined using illustrations). Define *portfolio value* as the ‘wealth’ of a portfolio using *value* as the price of good 1 and unity as the price of good 2. At the end of each stage, a subject was presented with his *value* and the *portfolio value* of both his stage-ending portfolio and his endowment (a subsidizer’s stage-ending portfolio value will be smaller than that of his endowment), and this data remained on-screen for the remainder of the session. Subjects were also presented with the *value lines* (i.e., budget lines) through their stage-ending portfolios. Subjects were given no instructions on how to use this information, but were informed they were simply for reference purposes. Budget lines could be toggled off and on as desired. See Figure 4 for a screen-shot of the user interface. Here, the human subject (Player 1) is a subsidizer, and the CSS-ZI robot is a subsidy receiver.

After running Treatments 1 through 3, there was little evidence to support CSS learning. A potential confound is that subjects were not trading in prices but rather portfolios; in hindsight the interpretation of one’s stage-ending marginal

Fig. 4 User Interface, Second Period of Play

rate of substitution as a price, crucial to CSS, was unnatural in this institution. **Treatment 4** was designed to focus more attention on this price and its implication for wealth redistribution. Subjects were instructed that when the relative value of goods between them was approximately equal, no more gains from trade were available. At such a time the current stage would be terminated immediately; thus, a stage could end before five minutes had expired if subjects implemented a near-Pareto optimal outcome. Subjects were informed that having a negative portfolio value (i.e., being a subsidizer) at the end of a stage is a signal that one has been outperformed in the exchange process, since subsidizers do not get the full termination value of the goods they have ‘sold.’ This treatment was a robustness check against the artificiality of subsidization constraint presentation in the others, and the results were reasonably similar.

4 Results

Several results emerged from the experiment that can be used to address the paper’s motivating questions. Subjects were generally quite adept at exhausting gains from trade, typically foregoing at most a penny or two in mutually advantageous gains after each stage was completed. In Treatments 1 and 4 it is apparent that some subjects learned to pin the CSS-ZI robots to their subsidization constraints and implement near-competitive outcomes in the latter stages. Evidence of adaptive behavior in the other two treatments is also present, although it is more difficult to disentangle due to feedback between subjects. However, most subjects (almost 70%) did not appear to target certain aspiration levels when making decisions, and the bulk of this group can be characterized as stochastic ZI or stochastic ZI with a propensity to tremble more frequently for small proposed changes in utility than for large ones (ZI+).

4.1 Exhausting gains from trade

Subjects typically reached a near-Pareto optimal allocation in each stage. One measure of convergence is the ratio of the distance between an allocation $\mathbf{x}$ and some ‘nearby’ point γ on the Pareto set to the distance between γ and ω :

$$\frac{\sqrt{(x_1 - \gamma_1)'(x_1 - \gamma_1)}}{\sqrt{(\omega_1 - \gamma_1)'(\omega_1 - \gamma_1)}}$$

A reasonable choice for γ is the intersection of the Pareto set with the line defined by $\mathbf{x}$ and (minus) the mean MRS at $\mathbf{x}$. The data confirm that most subjects seldom left any mutually beneficial gains from exchange unexploited. The median distance of final allocations from the Pareto set was quite flat across stages, below 0.01 in Treatment 1 and around 0.02 in the other treatments. A distance of 0.01 typically corresponds to 0 or 1 cents of lost gains from exchange, so median convergence was tight throughout. The mean distance was substantially higher in the first stage across treatments but dropped very close to the median by the tenth stage in all but Treatment 2, where two (of eight) economies skewed the mean upward.

4.2 Aggregate evidence of learning in Treatments 1 and 4

Recall that in ZI simulations of the Treatment 1 and 4 parametrization, only 5.1% of the time did Player 1 earn at least the competitive payment (\$1.24). In the first stage of the Treatment 1 and 4 experiments, the results were quite similar: Only one subject in each treatment earned at least this amount, corresponding to 5.6% and 4.6% of the subjects, respectively. However, by the tenth stage, three subjects in Treatment 1 (16.7%) and six subjects in Treatment 4 (27%) earned at least the competitive amount. Using the ZI simulations as the true distribution of outcomes under ZI play, the probability that three (six) ZI subjects would earn at least the competitive utility in any given period is 0.061 (0.001).

Figures 5-8 display the final allocations in stages 1 and 10 of Treatments 1 and 4. Defining distance as above, a 0.06-radius ring has been drawn about the competitive equilibrium to represent ‘near-competitive’ allocations. The minimum utility associated with this ring is \$1.10. Two of 18 (11.1%) subjects in Treatment 1 and 1 of 18 (4.6%) subject in Treatment 4 earn at least this much utility in the first stage, relative to 15.2% of the ZI simulations. By the tenth stage, these numbers grow to six (33.3%) and seven (31.8%), respectively. The probability that these proportions of ZI agents would earn at least \$1.10 in a period is less than 5%. It appears to be a reasonable hypothesis that ZI behavior was prevalent in the first stage, but by the tenth stage some subjects had become more discriminating. To test this hypothesis, it will be necessary to study individual decision-making.

4.3 Behavioral identification

In light of the ZI hypothesis, it is interesting to consider the percentage of point-increasing proposals (*gains*) and point-reducing proposals (*losses*) accepted by

Fig. 5 Treatment 1, stage-ending allocations in 1st stage

Fig. 6 Treatment 1, stage-ending allocations in 10th stage

Fig. 7 Treatment 4, stage-ending allocations in 1st stage

Fig. 8 Treatment 4, stage-ending allocations in 10th stage

subjects. Averaging across treatments, the median proportion of gains (losses) accepted was 94.4% (3.0%) (the individual treatments are similar except that the median proportion of losses accepted was substantially higher in Treatment 4, at 5.4%). Many subjects did accept far fewer gains and many more losses than the median. However, since 42% of subjects accepted at least 97.0% of their proposed gains (and in general these subjects accepted few losses), it may well be the case that near-ZI behavior is fairly common.

The following behavioral identification strategy was adopted. For each subject a preliminary model of decision-making was estimated, within which was nested a stochastic ZI model. Next, a structured attempt was made to reduce this baseline

model to its relevant explanatory variables. Then, if the estimated model suggested more sophisticated behavior than ZI, an attempt was made to determine if CSS provided any explanatory power, or if some alternative adaptive behavior might be present. Subjects were grouped into four behavioral types: Randomizers, ZI, ZI+, and aspirants. Randomizers exhibited no statistically significant difference in their reaction to gains and losses. ZI subjects exhibited a statistically significant preference for gains relative to losses, but did not significantly react to the magnitude of gains and losses or to proposed terms of trade (which will be defined shortly). ZI+ subjects were more likely to reject (accept) small gains (losses) than large ones, but otherwise behaved as ZI subjects. Aspirants attempted to secure advantageous terms of exchange by making the probability of accepting a proposal a function of the terms of trade, at least for some substantial length of time.

It is necessary to introduce some notation. Suppose subject i faces a total of $T + 1$ decisions during the experiment. Let $\mathbf{x}_{it}$ be the current portfolio of goods possessed by subject i when facing decision $t \in \{0, 1, 2, \dots, T\}$, and $\mathbf{r}_{it}$ be the portfolio proposed to i at t . Let z_{it} be the proposed change in points faced by subject i at t (i.e., $z_{it} \equiv u_i(\mathbf{r}_{it}) - u_i(\mathbf{x}_{it})$). Two indicator variables, g_{it} and ℓ_{it} , reflect when z_{it} is a gain or loss, respectively. That is, $g_{it} = 1$ if $z_{it} > 0$, zero otherwise, and $\ell_{it} = 1$ if $z_{it} < 0$, zero otherwise.

Another consideration pertains only to Treatments 1-3. Recall that in Treatment 4 a stage was terminated prior to 5 minutes if the subjects implemented a near-Pareto optimal allocation. Specifically, a stage was terminated if the difference in MRS between the two subjects was less than 0.40. Across subjects, 92% of the stages in Treatment 4 were terminated by this condition. The mean distance from Pareto optimality for terminated stages was 0.019, the maximum distance was 0.045. If the same condition had been applied in Treatments 1-3, 92% of the stages also would have been terminated. The mean proportion of decisions recorded after this (fictitious) stopping condition in Treatments 1-3 was nearly two-thirds. It seems plausible that the probability of accepting or rejecting a proposal after reaching a near Pareto-optimum could be quite different than earlier in the stage. Of course, subjects did not know explicitly they had reached a near-Pareto optimum, but it was generally the case that very few trades took place afterward, so subjects may have adapted to the change in pace of exchange. While it is potentially worrisome that subjects in Treatments 1-3 converged to the Pareto set so quickly (perhaps subjects became bored and stopped trying to do well), decision-making prior to near-Pareto optimality looks reasonably consistent with decision-making in Treatment 4. Therefore, let ξ_{it} be an indicator variable for subject i at decision t which equals 1 if the near-Pareto optimality condition has been reached, 0 if it hasn't. Occasionally the complementary indicator $\bar{\xi}_{it}$ will be referenced, which equals 1 if the near-Pareto optimality condition has not been reached, 0 otherwise.

Finally, subjects might reveal a preference for favorable terms of exchange. For example, it is plausible, perhaps expected, that in comparing two proposals for identical *gains*, a subject would prefer a proposal to trade 1 unit of good 1 for 5 units of good 2 to a proposal to trade 5 units of good 1 for 15 units of good 2. In the former case he would receive an implicit price of 5 for good 1, whereas in the latter he would only receive a price of 3. Let θ_{it} be the proposed terms of trade, in radians, faced by subject i at decision t . Specifically, $\theta_{it} = \arctan\left(\left|\frac{\mathbf{r}_2 - \mathbf{x}_2}{\mathbf{r}_1 - \mathbf{x}_1}\right|\right)$,

where the subscripts in the fraction refer to the type of good; subject and time subscripts have been suppressed for expositional clarity.⁵ A high value for θ_{it} (which is bounded by the proposal grid to be between 0.01 and 1.58) represents good terms of trade if the subject is ‘buying’ good 2, and poor terms if he is ‘selling’ good 2. Thus θ_{it} will enter the econometric model multiplicatively with an indicator variable for either being a buyer or a seller of good 2. Let $b_{it} = 1$ if the proposed net change in good 2 is positive, 0 otherwise, and let s_{it} be the complementary seller indicator.

Let $\Delta V_i(\mathbf{x}_{it}, \mathbf{r}_{it})$ be the difference in expected value for subject i , in the role of Player 1, between accepting and rejecting the current proposal. A probit model will be estimated under the assumption that the following (differenced) value function for subject i at decision t is distributed normally with mean 0 and unit variance:

$$\begin{aligned} \Delta V_i(\mathbf{x}_{it}, \mathbf{r}_{it}) = & \beta_{1i} * z_{it} + \beta_{2i} * g_{it} + \beta_{3i} * \ell_{it} + \beta_{4i} * \xi_{it} + \beta_{5i} * g_{it} * \xi_{it} + \\ & \beta_{6i} * \ell_{it} * \xi_{it} + \beta_{7i} * b_{it} + \beta_{8i} * g_{it} * b_{it} + \beta_{9i} * \ell_{it} * b_{it} + \\ & \beta_{10i} * \xi_{it} * b_{it} + \beta_{11i} * g_{it} * \xi_{it} * b_{it} + \beta_{12i} * \ell_{it} * \xi_{it} * b_{it} + \\ & \beta_{13i} * g_{it} * \xi_{it} * b_{it} * \theta_{it} + \beta_{14i} + \varepsilon_i. \end{aligned} \quad (1)$$

If subject i is assigned the role of Player 2, then b_{it} is replaced by s_{it} in the thirteenth term of Equation 1. Most proposed gains, and virtually all Pareto-improving proposals, faced by Player 1 (2) involve the exchange of good 1 (2) for good 2 (1)⁶, so the assumption that subjects ignore terms of trade on the side of the market in which they rarely trade is reasonable. Many other interactions were suppressed to avoid over-specifying the model. In particular, it is assumed that terms of trade do not influence the likelihood of accepting non-gains. It should be noted that the estimated coefficients of $\hat{\beta}_{7i}$ through $\hat{\beta}_{12i}$ are not necessarily expected to be significant influences on behavior. However, it is possible that some subjects will react differently to proposals depending on whether they are buying or selling good 2. For example, if the proposed gains for a subject in the role of Player 1 are mostly proposals to buy good 2, the subject may be more inclined to accept a proposal to sell good 2 just to ‘shake things up.’ The operating hypothesis is that these estimated coefficients are jointly insignificant (and if so they are dropped), but in many instances this hypothesis is rejected at the 5% level of significance. It should be noted that subjects in Treatments 1-3 made an average of 744, 1,000, and 1,027 total decisions, respectively, vs. 262 in Treatment 4.

For a subject in the role of Player 1 (2) in any of the treatments to be classified as an *aspirant*, $\hat{\beta}_{13i}$ will be greater (less) than zero and statistically significant at the 5% level. This outcome is a reasonably robust indication that at least for some length of time during the experiment, the subject attempted to hold out for good terms of trade by rejecting poorer ones. A reasonable interpretation of such behavior is that the subject is attempting to exert market power, or take a strong bargaining position. A less strategic interpretation (closer to the CSS story) is that

⁵ The angular normalization is used since the terms of trade parameter to be estimated will enter linearly.

⁶ The intuition for this statement should be clear after referring back to Figure 4, where Player 1 is near his initial endowment. Almost any proposal to buy good 1 would result in a loss, and no proposal to buy good 1 can be Pareto-improving.

the subject simply expects the market to move in a more favorable direction relative to his current position, and abstracts from his own influence on that market. The test is robust because the impact of the magnitude of proposed gains has been controlled for, and there is no obvious incentive for a subject to reject good terms of exchange for strategic reasons (otherwise, a subject might attempt to demand good terms in one stage and bad terms in another, the latter action potentially masking the former). A subject only sees ‘market decisions,’ so if he rejects a proposed trade, the action of his partner is not revealed. Thus, if Player 1 rejects a proposal on good terms of trade, either Player 2 is not aware of the fact, or he learns that he accepted a proposal on poor terms of trade which Player 1 nevertheless rejected. This likely would signal that the subject is demanding even better terms.

For a ZI+ subject i , $\hat{\beta}_{13i}$ will either have the ‘wrong’ sign or be statistically insignificant at the 5% level (usually the latter), but $\hat{\beta}_{1i} > 0$ and significant, suggesting that the subject is more likely to reject (accept) small gains (losses) than large ones. For a ZI subject j , $\hat{\beta}_{13j}$ and $\hat{\beta}_{1j}$ will be individually and jointly insignificant at the 5% level. Further, $\hat{\beta}_{2i}$ or $\hat{\beta}_{3i}$ (or both) must be statistically significant, with the estimate of the former being larger than the latter. That is, a ZI subject is only concerned with a proposal’s designation as a *gain* or *loss* and not with the magnitude of the proposal or its implied terms of trade. The decision of a subject classified as ZI may vary significantly with Pareto optimality or his roll as buyer or seller. Finally, a *randomizer* exhibits no significant difference in his rate of acceptance for gains and losses.

It is important to emphasize that the model described by Equation 1 is not intended to pin down specific behavior, but simply to sort out which subjects appear to consider the magnitude of proposals or terms of exchange when making their decisions. It is likely the case, for example, that a model which included the terms $z_{it} * g_{it}$ and $z_{it} * \ell_{it}$ would fit most ZI+ and aspirants better than Equation 1. However, doing so would not permit the inclusion of both g_{it} and ℓ_{it} individually due to collinearity, thus making it more difficult test for ZI behavior (as either gains or losses would necessarily be sucked up by the constant term along with proposals for zero-point changes). Of course, a probit model which includes only a constant, a gain indicator, and loss indicator is mis-specified, as the model requires the tails of the cdf to go to zero and one. Thus, we settle for a ZI model that has a non-zero but insignificant estimated coefficient on z_{it} .

The following estimation strategy was adopted. The probit model specified in Equation 1 was attempted to be fit for each subject. Despite the outright elimination of many interaction terms, it is the case for most subjects that several terms were collinear with the others due to lack of variation in the data. For terms which predict a subset of observations perfectly (e.g., subject i accepted 6 of 6 proposed gains as a seller of good 2 prior to reaching a near-Pareto optimum), these observations were excluded from the regression, and the collinear terms suppressed. Once the model was estimated, the joint significance of the estimated coefficients for the included variables from β_{7i} through β_{12i} was tested at the 5% level by a likelihood ratio test. The largest subset of jointly insignificant variables was eliminated from the model, and the model was re-estimated, incorporating any variables that had been initially suppressed which are not collinear with the remaining variables.

This process was repeated until all possible terms between the seventh and twelfth were eliminated. The last estimation typically had strong explanatory power.

Next, the hypothesis $\hat{\beta}_{13i} = 0$ was tested at the 5% level. If significant, the subject was classified as an aspirant, and three new variables were added to the model to determine if the subject increased his demanded terms of trade over time. One was the indicator $g_{it} * \bar{\xi}_{it} * b_{it} * css_{it}$, where $css_{it} = 1$ if proposal t passes the CSS constraint, 0 otherwise. For subjects in the role of Player 1, the other two variables were $g_{it} * \bar{\xi}_{it} * b_{it} * s_{it}$ and $g_{it} * \bar{\xi}_{it} * b_{it} * s_{it} * \rho_{it}$. Here $s_{it} = 1$ if $\theta_{it} < \pi/4$, 0 otherwise, and $\rho_{it} = 1$ if observation t is in stage 6-10, 0 otherwise. If the estimated coefficient on the CSS term is positive and significant, it suggests the subject attempted to increase his terms of trade over time in a manner not inconsistent with CSS. If the estimated coefficient for $g_{it} * \bar{\xi}_{it} * b_{it} * s_{it} * \rho_{it} < 0$, it suggests the subject rejected terms of trade worse than 1 unit of good 2 for 1 unit of good 1 more frequently in the second half of the experiment than the first. For subjects in the roll of Player 2, the condition in the second and third variables is changed from $\theta_{it} < \pi/4$ to $\theta_{it} > \pi/4$. These additional terms, while crude, give some indication as to whether a subject attempted to increase his terms of trade throughout the experiment.

If the null hypothesis $\hat{\beta}_{13i} = 0$ was rejected, then a significance test was applied to the estimate for β_{1i} . If significant at the 5% level, the subject is classified as ZI+. If not, and the estimated coefficients for β_{1i} and β_{13i} are not jointly significant either, the subject is classified as ZI if there is a significant estimated difference in the probability of rejecting gains and losses, or a randomizer otherwise.

Two other points are noted. First, any subject who rejected at most two gains prior to reaching a near-Pareto optimum was automatically classified as ZI, as there was insufficient variation in the data to test for alternative behavior. Second, some subjects exhibited a marked change in the proportion of gains and/or losses accepted at some point during the experiment. If the evidence was strong that a change in behavior took place (for example, a subject who accepted 98% of his proposed gains in the first three stages but only 70% of his proposed gains in the final seven stages), an indicator variable for gains and/or losses in the relevant periods was included before attempting to eliminate estimated coefficients for β_{7i} through β_{12i} .

4.4 Estimation results

A total of 41% of the subjects were classified as ZI, 24% were classified as ZI+, 31% as aspirants, and 4% as randomizers.⁷ Of the aspirants, 42% were found to have attempted to improve their terms of trade at some point during the experiment. For only 3 subjects did the CSS subsidization constraint offer significant explanatory power. One-third of the subjects classified as ZI and ZI+ in Treatments 1-3 significantly adjusted their behavior after a near-Pareto optimal allocation had

⁷ Results of the individual regressions are available at http://ist.caltech.edu/~scrockett/LCE_0601_results.pdf. The raw data is available at http://ist.caltech.edu/~scrockett/LCE_0601_data.xls.

Table 1 Proportion of Subject Types, By Treatment

	T1	T2	T3	T4
Randomizer	0%	6%	0%	9%
ZI	50%	38%	50%	27%
ZI+	11%	19%	22%	41%
Aspirant	39%	38%	28%	23%

been implemented; all aspirants did so. These adjustments could work in either direction, as some subjects became more likely to accept proposed gains and reject proposed losses after the condition was reached, while others were more likely to reject gains and accept losses.

Table 1 displays subject classifications by treatment, where there are two noticeable differences. First, the proportion of subjects classified as ZI in Treatment 4 is smaller than the others, at 27% compared to 50%, 38%, and 50% in Treatments 1, 2, and 3, respectively. This outcome is perhaps expected, as subjects were instructed explicitly to consider their status as CSS subsidizers. It is also possible that the large number of post-Pareto optimal proposals typically faced by subjects in Treatments 1-3 encouraged some subjects to become less discriminating with respect to gains and losses in general.

The other noticeable difference in classification across treatments is that there are more ZI+ and fewer aspirants in Treatment 4 relative to the others. This is a bit surprising, as subjects in Treatment 4 were essentially instructed to consider being an aspirant type. However, if the aspirant criterion were bumped down to a 10% level of statistical significance from 5%, three subjects would change classification from ZI+ to aspirant in Treatment 4, and no classifications would change in Treatments 1-3. This would make the proportions of ZI+ and aspirants much more similar across treatments (27% ZI+ and 36% aspirant in Treatment 4). The reason the aspirant classification criterion was not made less restrictive is because it seemed be the case with these particular subjects that the magnitude of point changes was a more economically and statistically significant influence on behavior than terms of trade. In any case, the classification of subjects is fairly similar across treatments, with the exception that there are appreciably fewer ZI subjects in Treatment 4.

Aspirants tended to earn more points than other subjects, and randomizers earn the least. The various parameterizations and rolls make it difficult to compare the performance of behavioral types across treatments. One way is to treat the mean simulated ZI payment for each treatment/player pair as benchmarks, and compute the percentage deviation of each subject's payment from the relevant benchmark. The mean simulated ZI (10-stage) payment for Player 1 in the first parametrization is \$8.80, and for Player 2 it is \$11.30. For Player 1 in Treatment 3, the ZI simulated mean payment is \$9.50, and for Player 2 the mean is \$10.50. The change in earnings as a proportion of the appropriate mean benchmark was calculated for each subject. Not surprisingly, randomizers fared the worst, earning 18.6% less on average than their simulated ZI counterparts. ZI subjects earned one-half of one percent more on average than their simulated ZI counterparts, so their earnings are nearly indistinguishable. Aspirants earned 15.8% more on average than the simulated ZI benchmarks. The big surprise was the earnings of

ZI+ subjects. On average, a ZI+ subject earned 5.3% less than his simulated ZI counterpart. The median ZI+ subject fared a little better, earning 0.6% less, but it is nevertheless surprising that these subjects did not earn more.

There are two partial explanations for why ZI+ subjects did not perform better. First, obtaining good terms of trade is a necessary condition for earning a lot of points, and it doesn't matter if those good terms are obtained in a couple of big steps worth many points each, or a series of small point-increasing proposals. In fact, across subjects in the role of Player 1 in Treatments 1, 2, and 4 (which account for almost 60% of all subjects), the correlation coefficient between proposed gains and proposed terms of trade is only 0.056. However, while this may help explain why ZI+ subjects don't do much better than ZI subjects, it does not explain why they actually do worse. It turns out that ZI+ subjects, along with aspirants, are much more likely to accept losses than ZI subjects.

Across all treatments, the median ZI subject accepted 1.3% of his proposed losses, whereas the median ZI+ subject accepted 6.6%, the median aspirant accepted 6.0%, and the median randomizer accepted 50%. In fact, while only 7% of the ZI subjects accepted at least 10% of their proposed losses, about 40% of both ZI+ and aspirants did so. While 63% of the ZI subjects accepted less than 2% of their proposed losses, only 18% of ZI+ subjects and 13% of aspirants accepted as few. Further, the proportion of subjects who accepted more than 10% of their proposed losses does not decrease much over time. While no ZI subjects accepted such a high proportion of losses in the final five stages, 35% of the ZI+ and aspirant types still accepted at least 10% of their proposed losses. Quite a few subjects did decrease the proportion of losses they accepted over time. Eighty percent of ZI subjects accepted fewer than 2% of their losses in the final five stages, whereas 35% of ZI+ and aspirants accepted as few.

The impact on earnings of the fairly large proportion of losses accepted by ZI+ and aspirant types on earnings is not at all straightforward, however. For randomizers and ZI subjects, the correlation coefficient between the proportion of losses accepted and earnings is large and negative (-0.27 and -0.40, respectively). For aspirants, though, the same correlation coefficient is -0.02, and for ZI+ subjects it is actually positive, at 0.05. Thus, the ZI+ subjects who accept a larger proportion of losses actually tend to earn a little more than the others.

It is likely the case that ZI+ and aspirant types generally accept losses to get a system "unstuck" from a Pareto optimum (after all, they do not know it is Pareto optimal, and may not even have any familiarity with the concept), or to experiment and make certain there are no hidden benefits to accepting losses (e.g., perhaps one's trading partner becomes more cooperative if he is somehow made aware that you accepted a loss, or if his decisions are conditioned on one's own rate of proposal acceptance). Consider Subject 27 in Treatment 2, who earned the highest payoff of all 74 subjects (\$17.19) while accepting 11% of his proposed losses, 24% over the final two stages. Subject 27 plays the roll of Player 2, and is matched with a human Player 1 who is classified as ZI+ (Player 1 was actually deterministically ZI until the final two stages). In Figure 9 we see that Player 2 has managed to pin Player 1 near his original indifference curve even while accepting many proposed losses; most of those losses could not have been accepted (of course, that means accepting them did not signal anything, because the other player would not have observed the signal). Similarly, in Figure 10 is presented an aspirant Player

Fig. 9 Treatment 2, Subjects 19 and 27, Stage-ending allocations

Fig. 10 Treatment 3, Subjects 38 and 47, Stage-ending allocations

1 who accepted 10% of his proposed losses in the final stage and still managed to pin his trading partner near the boundary of his individually rational set.

Thus, some aspirants occasionally accept losses in an apparent attempt to signal or explore and nevertheless do quite well, even if the losses are occasionally transacted. However, it has already been demonstrated that the magnitude of gains and terms of trade are very weakly correlated, so ZI+ subjects who experiment with losses do not reap much by way of benefits from the experimentation (even if they somehow get their trading partner to become more cooperative, they still are not holding out for good terms of trade), but may occasionally be hit with significant losses. This is one possible explanation for why ZI+ subjects earn somewhat less than ZI subjects.

5 Conclusion

The evidence presented in this paper suggests that behavioral heterogeneity can be a substantial influence in markets, and theories of equilibrium coordination that encompass a range of behavior, from zero intelligence to sophisticated, could prove to be descriptively relevant. More than two-thirds of the subjects in this experiment do not appear to set aspiration levels to guide their decision-making, but rather are myopic in their response to each exchange opportunity. Crockett, Spear, and Sunder propose that agents aspire to subsidize less than before. An alternative aspiration might be based explicitly on best expected terms of exchange. But for the randomizers, ZI, and ZI+ types identified in this paper, there is apparently no goal related to stage-ending allocations that guides decision-making.

Further, more than half of the subjects who did set aspiration levels did not appear to adapt them to market conditions. Gains from trade were completely exhausted in 92% of all stages, typically less than half-way through each stage. Thus, subjects had sufficient time to at least attempt to reach aspiration levels

beyond those associated with ZI play. They had simple means to do so (rejecting proposals with inferior terms of trade). They had market power. And the value of signalling cooperation (or the cost of signalling aggressiveness) was muted by the fact that only market decisions were reported, not necessarily individual decisions. It seems that this should be a rich environment in which to set fluid aspiration levels, and most subjects did not.

These findings do not present an outright rejection of CSS or cast new doubt on general equilibrium theory, as the evidence comes exclusively from a highly stylized market with only two participants. For instance, an institution where exchange is denominated in prices rather than consumption bundles might better lend itself to consideration of subsidization constraints specifically, or setting and monitoring aspiration levels more generally. Further, even if the absence of aspiration levels for a substantial subset of agents is found to be robust across other institutions (a non-aspirant in a double auction might execute any available utility-improving market order rather than form an expectation of what the best transaction will be), it is plausible that in large markets competitive equilibrium still might be obtained. For example, perhaps even a small proportion of aspirants would Bertrand-compete to quickly transact with non-aspirants, which might lead to competitive equilibrium as a limiting outcome. While non-aspirants may clearly be exploited in a small market when matched with sophisticated participants (Figures 9 and 10 present striking examples), they are not *necessarily* doomed to failure in large markets.

While these questions may suggest directions for future theoretical research, the analysis of behavior in this paper is potentially useful in its own right. In posted offer experiments it has been found that some subjects attempt to realize aspiration levels when market power provides them with an explicit means to do so while others do not (see the survey Davis and Holt (1993), or Cason and Reynolds (2005) for more recent research). An analogous result obtains here, but in a setting with some features of a more competitive market (for example, individuals do not set terms of trade, but rather decide when the market is ‘favorable enough to jump’). This support for aspirant/non-aspirant heterogeneity in a market with some competitive features may be useful in identifying behavior in critically important but more complicated institutions. The space of available actions in a double auction, for example, is huge (a participant must decide when to submit various limit orders and at what price and quantity). To estimate a serious econometric model of individual behavior in such an institution will no doubt require substantial prior structure on behavior.⁸ If some non-aspirants populate a double auction, the state and strategy space used to estimate a model of their behavior will be a fraction of the potential state and strategy space. Hopefully this paper offers value in refining the questions we ask about market coordination in general, in addition to identifying behavior in one class of markets which may relate to other, more complicated ones.

⁸ Consider the parametric learning models of (Camerer, Ho and Chong, 2002) (McKelvey and Palfrey, 1995), or the relatively small state spaces which are exploited in the behavioral typing work by Houser, Keane and McCabe (2004) and El-Gamal and Grether (1995), and then imagine how the strategy and state space of a double auction dwarf these environments in comparison.

Appendix: Subject Payment and Type Information

Subject ID	Type	Treatment, Roll	Partner ID	Partner Type	% loss accept	Pay	Mean ZI-sim pay for Treatment, Roll	% Dev. From ZI
31	Random	T2, Player 2	23		0.50	6.82	11.30	-0.40
59	Random	T4, Player 1	n/a	ZI-CSS robot	0.46	8.24	8.80	-0.06
74	Random	T4, Player 1	n/a	ZI-CSS robot	0.53	7.95	8.80	-0.10
3	ZI	T1, Player 1	n/a	ZI-CSS robot	0.03	8.28	8.80	-0.06
5	ZI	T1, Player 1	n/a	ZI-CSS robot	0.03	9.39	8.80	0.07
6	ZI	T1, Player 1	n/a	ZI-CSS robot	0.04	9.36	8.80	0.06
7	ZI	T1, Player 1	n/a	ZI-CSS robot	0.01	9.67	8.80	0.10
8	ZI	T1, Player 1	n/a	ZI-CSS robot	0.05	8.79	8.80	0.00
10	ZI	T1, Player 1	n/a	ZI-CSS robot	0.02	8.82	8.80	0.00
11	ZI	T1, Player 1	n/a	ZI-CSS robot	0.00	9.13	8.80	0.04
12	ZI	T1, Player 1	n/a	ZI-CSS robot	0.00	8.87	8.80	0.01
18	ZI	T1, Player 1	n/a	ZI-CSS robot	0.06	8.52	8.80	-0.03
22	ZI	T2, Player 1	30	ZI+	0.01	8.94	8.80	0.02
24	ZI	T2, Player 1	32	ZI	0.01	7.86	8.80	-0.11
26	ZI	T2, Player 1	34	ZI	0.01	10.17	8.80	0.16
28	ZI	T2, Player 2	20	Aspirant	0.01	10.51	11.30	-0.07
32	ZI	T2, Player 2	24	ZI	0.01	12.34	11.30	0.09
34	ZI	T2, Player 2	26	ZI	0.01	10.20	11.30	-0.10
35	ZI	T3, Player 1	44	ZI	0.15	6.99	9.50	-0.26
36	ZI	T3, Player 1	45	ZI+	0.00	9.72	9.50	0.02
37	ZI	T3, Player 1	46	Aspirant	0.00	9.56	9.50	0.01
39	ZI	T3, Player 1	48	ZI	0.02	9.67	9.50	0.02
42	ZI	T3, Player 1	51	Aspirant	0.01	9.32	9.50	-0.02
43	ZI	T3, Player 1	52	ZI+	0.02	9.29	9.50	-0.02
44	ZI	T3, Player 2	35	ZI	0.00	13.93	10.40	0.34
47	ZI	T3, Player 2	38	Aspirant	0.01	8.28	10.40	-0.20
48	ZI	T3, Player 2	39	ZI	0.02	10.43	10.40	0.00
54	ZI	T4, Player 1	n/a	ZI-CSS robot	0.04	9.32	8.80	0.06
55	ZI	T4, Player 1	n/a	ZI-CSS robot	0.23	8.15	8.80	-0.07
57	ZI	T4, Player 1	n/a	ZI-CSS robot	0.03	8.97	8.80	0.02
60	ZI	T4, Player 1	n/a	ZI-CSS robot	0.03	8.40	8.80	-0.05
62	ZI	T4, Player 1	n/a	ZI-CSS robot	0.00	10.40	8.80	0.18
69	ZI	T4, Player 1	n/a	ZI-CSS robot	0.03	8.37	8.80	-0.05

Table 2 Randomizers and ZI

Subject ID	Type	Treatment, Roll	Partner ID	Partner Type	% loss accept	Pay	Mean ZI-sim pay for Treatment, Roll	% Dev. From ZI
2	ZI+	T1, Player 1	n/a	ZI-CSS robot	0.03	8.75	8.80	-0.01
9	ZI+	T1, Player 1	n/a	ZI-CSS robot	0.23	7.26	8.80	-0.18
19	ZI+	T2, Player 1	27	Aspirant	0.03	4.31	8.80	-0.51
29	ZI+	T2, Player 2	21	Aspirant	0.09	8.26	11.30	-0.27
30	ZI+	T2, Player 2	22	ZI	0.02	11.26	11.30	0.00
40	ZI+	T3, Player 1	49	Aspirant	0.24	9.53	9.50	0.00
45	ZI+	T3, Player 2	36	ZI+	0.02	10.31	10.40	-0.01
50	ZI+	T3, Player 2	41	Aspirant	0.07	9.61	10.40	-0.08
52	ZI+	T3, Player 2	43	ZI	0.00	10.71	10.40	0.03
53	ZI+	T4, Player 1	n/a	ZI-CSS robot	0.50	9.12	8.80	0.04
58	ZI+	T4, Player 1	n/a	ZI-CSS robot	0.13	8.62	8.80	-0.02
63	ZI+	T4, Player 1	n/a	ZI-CSS robot	0.10	8.83	8.80	0.00
65	ZI+	T4, Player 1	n/a	ZI-CSS robot	0.01	9.31	8.80	0.06
66	ZI+	T4, Player 1	n/a	ZI-CSS robot	0.19	8.28	8.80	-0.06
67	ZI+	T4, Player 1	n/a	ZI-CSS robot	0.03	9.10	8.80	0.03
71	ZI+	T4, Player 1	n/a	ZI-CSS robot	0.07	8.59	8.80	-0.02
72	ZI+	T4, Player 1	n/a	ZI-CSS robot	0.02	9.54	8.80	0.08
1	Aspirant	T1, Player 1	n/a		0.07	8.74	8.80	-0.01
4	Aspirant	T1, Player 1	n/a		0.01	9.90	8.80	0.13
13	Aspirant	T1, Player 1	n/a		0.14	11.29	8.80	0.28
14	Aspirant	T1, Player 1	n/a		0.19	9.75	8.80	0.11
15	Aspirant	T1, Player 1	n/a		0.03	10.29	8.80	0.17
16	Aspirant	T1, Player 1	n/a		0.08	9.65	8.80	0.10
17	Aspirant	T1, Player 1	n/a		0.01	9.95	8.80	0.13
20	Aspirant	T2, Player1	28	ZI	0.02	10.08	8.80	0.15
21	Aspirant	T2, Player 1	29	ZI+	0.28	13.34	8.80	0.52
23	Aspirant	T2, Player 1	31	Random	0.01	12.05	8.80	0.37
25	Aspirant	T2, Player 1	33	Aspirant	0.05	9.85	8.80	0.12
27	Aspirant	T2, Player 2	19	ZI+	0.11	17.19	11.30	0.52
33	Aspirant	T2, Player 2	25	Aspirant	0.03	10.39	11.30	-0.08
38	Aspirant	T3, Player 1	47	ZI	0.04	12.71	9.50	0.34
41	Aspirant	T3, Player 1	50	ZI+	0.08	10.62	9.50	0.12
46	Aspirant	T3, Player 2	37	ZI	0.14	10.34	10.40	-0.01
49	Aspirant	T3, Player 2	40	ZI+	0.14	8.59	10.40	-0.17
51	Aspirant	T3, Player 2	42	ZI	0.05	10.88	10.40	0.05
56	Aspirant	T4, Player 1	n/a		0.30	10.67	8.80	0.21
61	Aspirant	T4, Player 1	n/a		0.00	11.67	8.80	0.33
64	Aspirant	T4, Player 1	n/a		0.37	9.45	8.80	0.07
68	Aspirant	T4, Player 1	n/a		0.02	10.11	8.80	0.15
70	Aspirant	T4, Player 1	n/a		0.46	9.25	8.80	0.05
73	Aspirant	T4, Player 1	n/a		0.03	10.21	8.80	0.16

Table 3 ZI+ and Aspirants

References

- Anderson, C. M., C. Plott, K.-I. Shimomura and S. Granat (2004). "Global instability in experimental general equilibrium: The scarf example." *Journal of Economic Theory*, **115**(2), 209–249.
- Camerer, C., T. Ho and J. Chong (2002). "Sophisticated experience-weighted attraction learning and strategic teaching in repeated games." *Journal of Economic Theory*, **104**, 137–188.
- Cason, T. and S. Reynolds (2005). "Bounded rationality in laboratory bargaining with asymmetric information." *Economic Theory*, **25**, 553–574.
- Crockett, S., S. Spear and S. Sunder (2006). "Learning competitive equilibrium." *Journal of Mathematical Economics*, **forthcoming**.
- Davis, D. and C. Holt (1993). *Experimental Economics*. Princeton University Press, Princeton, NJ.
- El-Gamal, M. and D. Grether (1995). "Are people bayesian? uncovering behavioral strategies." *Journal of the American Statistical Association*, **90**, 1137–1145.

-
- Gjerstad, S. (2004). "Market dynamics in edgeworth exchange." Working paper.
- Gode, D., S. Spear and S. Sunder (2004). "Convergence of double auctions to pareto optimal allocations in the edgeworth box." Working paper.
- Gode, D. and S. Sunder (1993). "Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality." *Journal of Political Economy*, **101**(1), 119–137.
- Houser, D., M. Keane and K. McCabe (2004). "Behavior in a dynamic decision problem: An analysis of experimental evidence using a bayesian type classification algorithm." *Econometrica*, **72**(3), 781–822.
- Lian, P. and C. Plott (1998). "General equilibrium, markets, macroeconomics, and money in a laboratory experimental environment." *Economic Theory*, **12**, 21–75.
- McKelvey, R. and T. Palfrey (1995). "Quantal response equilibria for normal-form games." *Games and Economic Behavior*, **7**, 6–38.
- Noussair, C., C. Plott and R. Reizman (1995). "An experimental investigation of the patterns of international trade." *The American Economic Review*, pp. 462–491.
- Smith, V. (1962). "An experimental study of competitive market behavior." *Journal of Political Economy*, **70**(2), 111–137.
- Williams, A., V. Smith, J. Ledyard and S. Gjerstad (2000). "Concurrent trading in two experimental markets with demand interdependence." *Economic Theory*, **16**, 511–528.