Price Dynamics in an Exchange Economy

Steven Gjerstad

Economics Department
Purdue University
403 W. State Street
West Lafayette, IN 47907-2056
gjerstad@purdue.edu

November 19, 2007

Abstract The pure exchange model is the foundation of the neoclassical theory of value, yet equilibrium predictions and models of price adjustment for this model remained untested prior to the experiment reported in this paper. With the exchange economy replicated several times, prices and allocations converge sharply to the competitive equilibrium in continuous double auction (CDA) trading. Convergence is evaluated by comparing the extent of price adjustment within each market replication (or trading period) to the extent of adjustment across trading periods: most observed price adjustment occurs within trading periods, so price adjustment data are evaluated with the Hahn process model (Hahn and Negishi [1962]), which is a disequilibrium model of within-period trades. Estimation demonstrates that the model is consistent with observed price paths within each period of the exchange economy. The model is augmented with an additional assumption – based on observations from this experiment – that the initial trade price in period $t+1$ is randomly drawn from the interval between the minimum and maximum trade prices in period t. The estimated within-period adjustment rule, combined with this across-period adjustment rule, generates price paths similar to data from an experiment session.

Keywords: Competitive equilibrium, disequilibrium dynamics, continuous double auction, experimental economics, exchange economy, Hahn process, neoclassical theory of value, tatonnement, unit root tests

JEL Classifications: C22, C92, D41, D44, D51

1 Introduction

The pure exchange model forms the foundation of the neoclassical theory of value. Assessment of its price and allocation predictions is important in microeconomic theory as well as in much applied economics that rests on a general equilibrium foundation. This paper reports a direct experimental test of the competitive equilibrium predictions and price dynamics in an exchange economy with strong income effects. Income effects pose a challenge for the price discovery process in a general equilibrium model. Resource ownership is specified in a general equilibrium model, so the wealth of an agent depends on prices, and wealth affects demand. Except in the case of quasi-linear preferences, a disequilibrium trade alters the wealth of each party to the trade from their wealth in equilibrium. This alters agents' demands, and hence alters the equilibrium price of the economy that remains after the disequilibrium trade. Consequently, equilibrium may be much more difficult to reach in a market in which agents have income effects than in the experimental markets previously studied. This is not merely a subtle issue in general equilibrium theory. This situation arises in international trade. The large recent increase in oil prices produces a large transfer of income from oil importers to exporters. The process of price adjustment in the presence of income effects of this sort tells us how prices respond to altered market conditions. The question addressed by the experiment in this paper is whether a simple general equilibrium economy with strong income effects can reach its competitive equilibrium, and if so, how that comes about. The test reported here demonstrates that prices and allocations do converge to the competitive equilibrium, but the process of convergence is subtler that in previous experiments that do not involve income effects.

In principle, an experimental test of this model is straightforward. Smith [1976, 1982] developed a theory of induced utility which applies directly to this problem. The experimenter specifies a utility function for each agent in an exchange economy. The value of an agent's induced utility function at his final allocation determines his payoff. Any of several exchange mechanisms could meditate trades. Due to its robust performance properties and its operational simplicity, the continuous double auction (CDA) is a natural choice. A partial test of this type – with induced utility functions for buyers and schedules of induced costs for sellers – is reported in Williams, Smith, Ledyard, and Gjerstad [2000]. Prices and allocations in their experiment converge, albeit with a slight price bias. The test reported in this paper combines induced utilities for sellers and buyers, and unlike previous general equilibrium experiments, such as Lian and Plott [1998] or Anderson, Plott, Shimomura, and Granat [2004], the experiment uses a commodity money so that the absolute price level is determined, and the market environment is tightly linked to the exchange economy. Moreover, unlike previous general equilibrium experiments, all agents in the experiment reported in this paper have strictly quasi-concave utility functions that exhibit strong income effects, so the experiment provides a strenuous test of competitive equilibrium predictions.

A competitive market environment is created with six sellers and six buyers, each with an induced utility function over two commodities. One commodity is specified as the numeraire: prices for the non-numeraire commodity are stated in terms of a number of units of the numeraire. Buyers initially hold only the numeraire commodity in their endowment, and sellers initially hold only the non-numeraire commodity in their endowment.[1] Each experiment session consists of either 12 or 15 identical replications of the exchange economy.[2] In each of these replications (or trading periods), the sellers post offer prices and the buyers post bid prices – denominated in units of the numeraire commodity – for one unit of the other commodity. Exchange occurs when a seller and a buyer agree on a number of units of the numeraire commodity that they will exchange for a unit of the other commodity. Each agent may (and typically does) participate in several trades in each trading period.

Trades strongly support competitive equilibrium price and allocation predictions in this experiment. By the final period in four sessions, allocations and prices were very near the competitive equilibrium allocation and price: the allocation differed from the competitive equilibrium allocation by less than 3% (measured as the ratio of the distance between the final allocations and the equilibrium allocations to the distance between the initial endowments and the equilibrium allocations). In addition, the difference between the average final period price and the equilibrium price was less than 3.7% in each session and averaged less than 3.1%. Earnings as a percentage of equilibrium earnings were closer yet: in the final period they averaged 99.8%, which is higher than the typical efficiency in a CDA experiment with unit costs for sellers and unit values for buyers.

In addition to confirmation of convergence to competitive equilibrium, these experiment sessions provide a unique opportunity to assess price dynamics in an exchange economy. In the CDA trades are negotiated directly between sellers and buyers, trades typically occur at a variety of prices even within a single trading period, and most trades occur out of equilibrium. These characteristics of the trading process, rather than hindering investigation of price dynamics, facilitate their evaluation.

In order to shed light on this adjustment process, market dynamics are evaluated with two price adjustment models. Across-period adjustment is evaluated with the price adjustment rule from the tatonnement model. Although the price adjustment rule from the tatonnement model is normally interpreted as an action by the 'auctioneer' in a Walrasian auction, it is nevertheless useful as a diagnostic tool for evaluation of changes to average price across trading periods, even

[1] The position of each subject in the experiment as either a buyer or a seller is of course a fiction: those labeled 'buyer' in the experiment because they purchase units of the non-numeraire commodity are also sellers of the numeraire commodity; sellers of the non-numeraire commodity can also be viewed as buyers of the numeraire commodity.

[2] The first three sessions – all conducted on a single day – provided evidence that convergence is much slower in this environment than in a market with induced costs and values, so the number of market replications was increased for the fourth session.

in a market that employs the CDA trading mechanism. Individual excess demand is a well-defined function of an agent's current allocation and the market price, and market excess demand is the sum of individual excess demands. Market excess demand at the average price in period t leads under the tatonnement adjustment rule to a predicted change to the average price in period $t+1$. Consequently, we can use the price adjustment rule from tatonnement as a diagnostic tool, and determine whether it describes period-to-period price adjustment well, even though the impetus to these price changes differs in the CDA from that in the Walrasian auction.

The tatonnement price adjustment rule is therefore used to evaluate convergence across market periods, with the predicted average price in period $t+1$ determined as $\bar{p}_{t+1} = \bar{p}_t + b\,Z(\bar{p}_t) + \epsilon_{t+1}$. Under the null hypothesis of no price change in response to excess demand ($b = 0$), the process is a random walk. The linear approximation to this adjustment equation, $\bar{p}_{t+1} = \bar{p}_t + \beta\,(p^* - \bar{p}_t) + \epsilon_{t+1}$, is examined with Dickey-Fuller unit root tests. These tests demonstrate that average price adjusts across periods toward equilibrium, as predicted.

Evaluation of period-to-period average price adjustment though neglects the substantial within-period price adjustment. In the "Hahn process" (Hahn and Negishi [1962]), the auctioneer announces a price as in tatonnement; traders respond by announcing excess demand. Unlike tatonnement though, in this disequilibrium exchange model trade takes place if there is some exchange compatible with announced excess demands. Then price adjusts in response to excess demand, as in the tatonnement model, and the process is repeated at the new price. Proceeding iteratively, the model predicts a path of prices within a trading period. The Hahn process model is used to evaluate experiment data by comparing predicted within-period price adjustment to observed within-period price adjustment from experiment sessions. As with tatonnement adjustment, Hahn process adjustment is a random walk under the null hypothesis of no adjustment, so unit root tests are applied to within-period trades as well. The Hahn process accounts for and adapts to disequilibrium trades, and converges to a Pareto optimal allocation in each period. Its predicted price path within a period t, augmented with a predicted initial price in period $t+1$, effectively describes the evolution of prices both within and across trading periods.

When agents have income effects, convergence is slower and subtler than in markets with induced supply and demand which have no income effects. In a market without income effects, the equilibrium price near the end of the period approximately equals the equilibrium price before trading begins. This simplifies the learning problem since the end-of-period trades condition agents' subsequent price expectations, so that prices quickly converge to the equilibrium. In a market with income effects though, when the price path differs from the equilibrium price, the equilibrium price by the end of the period may differ substantially from the equilibrium price at the start of a trading period. With replication of the market conditions across several trading periods though, this difficulty is overcome. Prices within a trading period typically adjust toward the competitive equilibrium price, so that the price path in each period better approximates the CE price. Eventually prices and allocations stabilize in the vicinity of the equilibrium price and allocation. Although price paths in markets with income effects have different characteristics than price paths in markets without income effects, both are consistent with the same adjustment process.

This paper progresses through the steps outlined above. Section 2 describes the economic environment induced in the experiment sessions. Section 3 describes the CDA mechanism in the exchange economy context. Section 4 evaluates experiment data for across-period convergence and also compares across-period price adjustment to within-period price adjustment. Section 5 describes the within-period price adjustment from the Hahn process model, augmented with a simple across-period price adjustment rule. Section 6 reports Hahn process model parameter estimates and a simulation of the adjustment process. Conclusions are drawn in Section 7.

2 Economic environment

At the beginning of each trading period, each buyer is endowed with 1800 units of the numeraire commodity (X), which is described to the buyer as currency. Buyers can use this numeraire commodity to purchase units of the commodity (Y). Each seller is endowed with eighteen units of the commodity, which can be sold individually to acquire units of the numeraire commodity. The payoff of agent i is determined from the value of the constant elasticity of substitution utility function $u_i(x, y) = c_i\,((a_i\,x)^{r_i} + (b_i\,y)^{r_i})^{1/r_i}$, where x is the final allocation of currency (the numeraire commodity X) and y is the final allocation of the commodity (Y) held by agent i at the end of a trading period.[3]

In each period of an experiment session, each of six buyers (agent type B) has the same utility function $u_B(x, y)$ and endowment $\omega_B = (1800, 0)$; each of six sellers (agent type S) has the same utility function $u_S(x, y)$ and endowment $\omega_S = (0, 18)$. Table 1 shows the CES utility function parameters and the endowments of sellers and buyers.

	c_i	a_i	b_i	r_i	ω_i
Sellers' parameters $(i = 2, 4, \ldots, 12)$	0.256	2.982	109.89	-1	(0, 18)
Buyers' parameters $(i = 1, 3, \ldots, 11)$	0.695	0.362	109.89	-1	(1800, 0)

Table 1: Sellers' and buyers' parameters.

With $\rho_i = r_i/(1 - r_i)$, the excess demand for Y of agent i is

[3] The parameter c_i does not affect the competitive equilibrium price or allocation, but it does affect the utility level attained by subject i, and is therefore a useful element of the utility inducement technique, since it permits the experimenter to rescale a subject's payoff with a single parameter.

$$Z_i^Y(p \mid \omega_i) = \frac{b_i^{\rho_i}(\omega_i^X + p\,\omega_i^Y)}{p\left(b_i^{\rho_i} + p^{\rho_i} a_i^{\rho_i}\right)} - \omega_i^Y . \tag{2.1}$$

For commodity Y, market excess demand $Z^Y(p \mid \omega) = \sum_{i=1}^{12} Z_i^Y(p \mid \omega_i)$ is the sum of individual excess demands.[4] When parameters from table 1 are substituted into equation (2.1) and the resulting individual excess demand functions are substituted into the market excess demand function, the equilibrium price $p^* \doteq 91$ is obtained as the solution to $Z^Y(p \mid \omega) = 0$. (In equilibrium, the market for the numeraire commodity (X) then clears by Walras' law.) Then $p^* = 91$ is substituted into equation (2.1) to determine the net trade of Y by agent i. Each seller has equilibrium allocation $(x_S, y_S) = (637, 11)$; each buyer has equilibrium allocation $(x_B, y_B) = (1163, 7)$. The equilibrium utility level for each agent is 189. With sellers identical to one another, and buyers identical to each other, the economic environment can be displayed in an Edgeworth diagram, as in figure 3.

In order to implement an exchange economy experiment, each subject is provided with detailed specifications of his objective and of the exchange institution. The next section describes the CDA institution in the context of an exchange economy and describes the representation of the utility inducement technique to subjects.

3 The continuous double auction (CDA)

In the CDA, any seller may submit an ask at any time during a trading period. An ask, which is the seller's current report of the fewest units of the numeraire commodity that he is willing to accept for a unit of the commodity, is entered in the area on the seller's screen display labeled "Enter Ask," as in figure 1. Similarly, a buyer's bid, which represents her current report of the most units of the numeraire commodity that she is willing to pay for a unit of the commodity, may be submitted at any time. An ask placed at or below the current high bid generates a trade at the bid price. A bid that meets or exceeds the current low ask yields a trade at the ask price. A seller may make any number of asks, and may trade any number of units that is consistent with his commodity endowment. Similarly, a buyer may make any number of bids, and may trade any number of units that is consistent with her endowment of the numeraire commodity. Several specific rules are implemented in the version of the CDA used in this experiment. Of these, the most important is the "spread reduction rule," which requires that each new ask is below the current low ask and each new bid exceeds the current high bid. A seller has the option to remove any ask that he has previously made, provided his request to remove the ask is received before it results in

[4] Dependence of excess demand on endowments ω_i or on interim allocations (x_i, y_i) is indicated because in the Hahn process, price adjusts after each trade in response to excess demand at the current allocation. In Section 5, which defines the Hahn process adjustment rule, the excess demand for Y is written $Z^Y(p \mid (x, y)) = \sum_{i=1}^{12} Z_i^Y(p \mid (x_i, y_i))$.

a trade. Each seller is permitted a single ask in the market queue at any time. Any new ask by a seller replaces his previous ask if he has one in the queue. Each ask is the unit price offered for one unit: multiple unit trades are not permitted. Analogous restrictions apply to each buyer's bids.

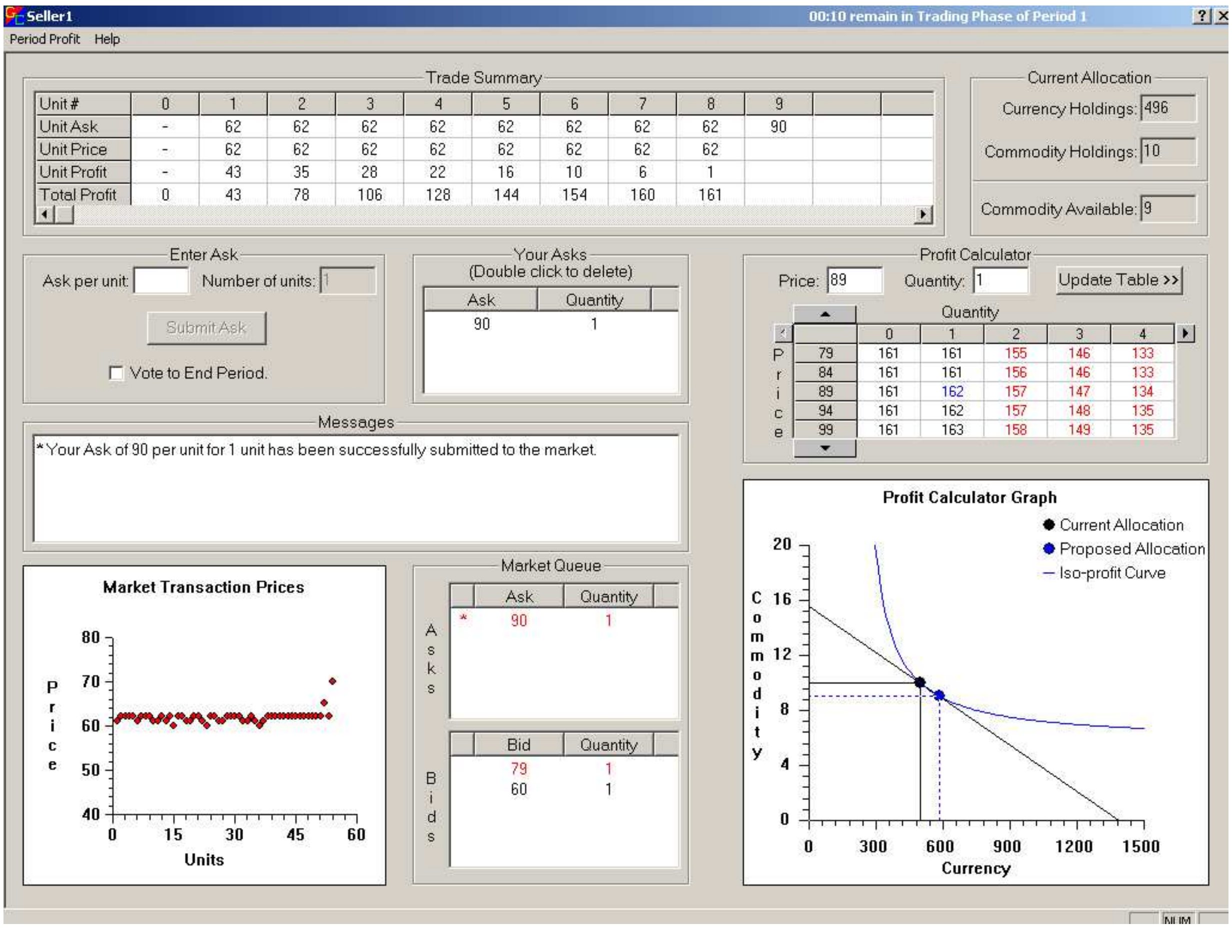

Figure 1: Seller screen with elements of market institution.

During each period, a queue on the seller's screen displays all current asks and bids (shown as the "Market Queue" in figure 1); each buyer's screen also displays both queues. When a seller successfully enters an ask into the ask queue, he receives a confirmation message in the "Messages" area of the screen display. This ask also appears in the "Unit Ask" row of the "Trade Summary" table, in the column that corresponds to the unit the seller has offered for sale. Similarly, a buyer receives a confirmation message when she enters a bid into the bid queue, and also sees an update to the appropriate cell in her Trade Summary table. When a seller and buyer complete a trade, they both receive a confirmation message, "Unit Price," "Unit Profit," and "Total Profit" figures are recorded in their Trade Summary tables, and the price appears in the "Market Transaction Prices" graph, which displays all trade prices from the current period. The trading phase of each period lasts 180 seconds; a clock at the top of each agent's screen shows the time remaining in the current phase. Each of these elements of the market institution appears on a seller's trading screen, as in figure 1. Buyers have a similar trading screen.

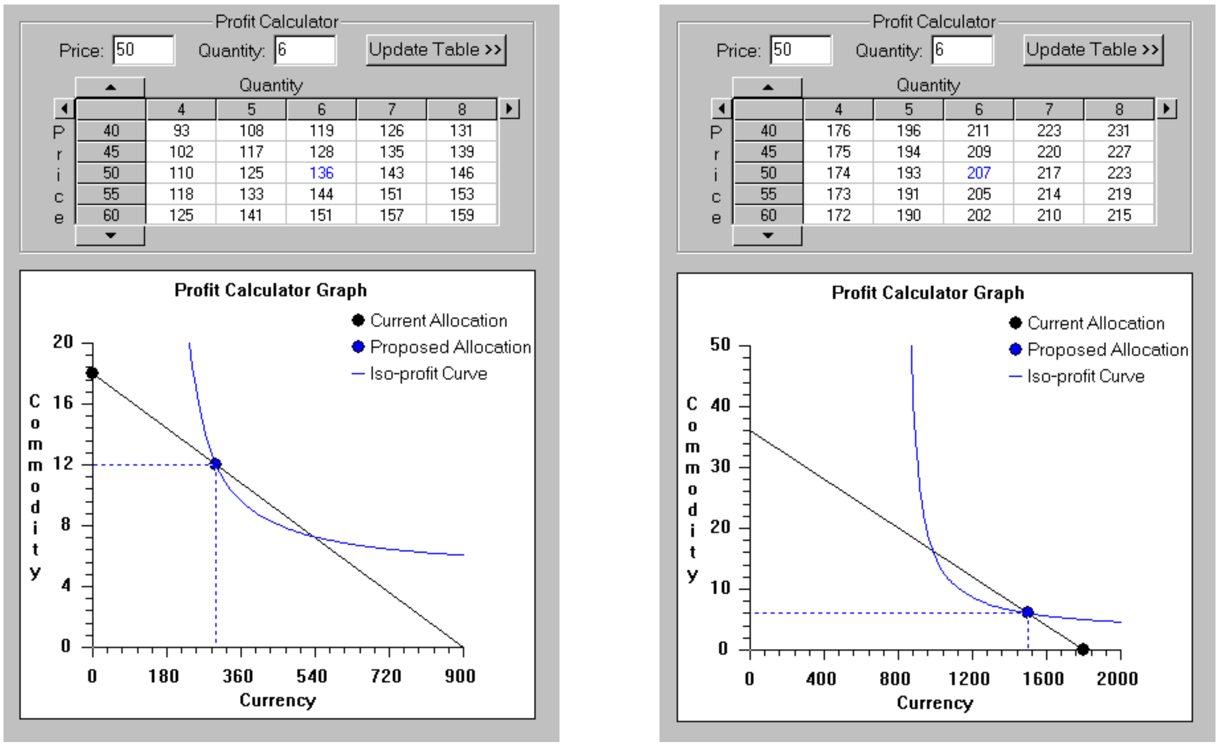

Figure 2: Profit calculator table and graph for seller (left) and buyer (right).

Profit (or utility) representation

Each market period is separated into three phases. During the preview phase, which lasts 60 seconds, and during the 180 second trading phase, a seller is able to enter "Price" and "Quantity" into a "Profit Calculator," which appears on the lower right hand side of the seller screen in figure 1. (The final phase is the 30 second review phase, when sellers and buyers have an opportunity to examine the results of the trading phase.) This Profit Calculator displays both a tabular and a graphical representation of the utility level that would result from a proposed trade. Examples of the Profit Calculator are shown for a seller and for a buyer in figure 2, starting from the seller's and the buyer's initial endowments. When a subject enters data into the Price and Quantity boxes in this calculator and clicks "Update Table," the profit (or utility) level is displayed in the center of the table for the allocation that would result from the proposed exchange. Profit levels are also displayed in the table for prices above and below the proposed price, and for quantities above and below the proposed quantity. In addition, the graph represents the "Current Allocation," the "Proposed Allocation" which would result from the proposed Price and Quantity, and the "Iso-profit Curve" (or indifference curve) that passes through the Proposed Allocation.

Instructions

The CDA mechanism and the representation of profit (or utility) levels are explained to subjects in a detailed interactive instruction set. Instructions describe each element of the seller's (or buyer's) trading screen independently, and then describe how elements relate to one another. There are points in the instruction set at which the seller or buyer is prompted for inputs, and there are eight interactive questions that must be answered correctly in order to proceed through the instructions. The sequence of steps through the instructions is outlined in Appendix A.

Experiment sessions and experience

Experiment sessions York1, York2, and York3 were conducted at the University of York in the U.K. on June 7, 2001. These three sessions had a total of thirty-six student subjects drawn from a wide range of backgrounds, including the physical and social sciences, the arts, and the humanities. Session Ariz1 was conducted at the University of Arizona on Nov. 16, 2001 with twelve undergraduate students recruited from across all majors.[5] Subjects in all exchange economy sessions had prior CDA experience in a session with a list of unit costs for each seller and a list of unit values for each buyer, rather than a utility function. Experience in a CDA with induced costs and values simplified subjects' learning tasks, since they were already familiar with the CDA market institution when they participated in the induced utility experiment. From the pool of 120 subjects at York who had experience in an induced cost and value CDA in the two weeks prior to sessions York1, York2, and York3, for the induced utility sessions we did not attempt to recruit the 20 subjects whose earnings were lowest as a fraction of equilibrium earnings in the induced cost and value sessions. A similar criterion was applied for session Ariz1.

[5] In a fifth session, prices continued to decline between periods 10 and 15, after they had already fallen to the competitive equilibrium price. During these six periods, 22 units were sold by the sellers at a loss. Misunderstanding of the utility inducement technique, particularly by three sellers, led to their poor performance and to a large deviation from equilibrium. Since the supply exhibited by these sellers substantially exceeded the induced supply, analysis of the market requires a decision error variable in the regressions. This analysis though complicates the basic adjustment model, provides no additional insight into price dynamics, and doesn't alter the basic conclusions of the paper, so the fifth session is not evaluated in this paper.

It is not the failure to converge though that limits the usefulness of the data from the fifth session. An adjustment process could fail to converge, even though no subject trades at a loss. If price adjustment includes both lagged price changes (with positive coefficients) and the usual adjustment term from excess demand, then the price path may repeatedly overshoot the equilibrium and fail to converge. This would be an interesting outcome well worth analysis. The failure to converge in the fifth session though did not result from an unanticipated adjustment dynamic like this, but rather from an apparent misunderstanding of the utility inducement technique by three sellers.

4 Convergence

This section assesses convergence of prices and allocations in four experiment sessions. Section 4.1 assesses convergence of individual allocations to the competitive equilibrium allocation. Section 4.2 applies the tatonnement price adjustment rule to evaluate across-period convergence of mean prices. Section 4.3, which compares the extent of across-period price convergence to the extent of within-period price convergence, demonstrates that most price adjustment occurs within trading periods. This motivates the analysis of within-period adjustment in Sections 5 and 6.

4.1 Convergence of allocations across periods

Prices converge reliably in the exchange economy experiment sessions (with the exception of the session in which several subjects apparently misunderstood the utility inducement technique). Figure 3 (a) shows the per capita interim allocations after each trade during three trading periods from session York1. Figure 3 (b) depicts final per capita allocations from each period in York1. Figure 4 shows, in detailed views, final allocations by period for each session. Note that the diagrams in figure 4 show only 2.6% of the area of the Edgeworth diagrams in figure 3, and the diagram for session York2 shows less than 0.25% of the area. These diagrams suggest a significant level of convergence across periods in three of four sessions, and almost immediate convergence in the first period of the session shown in figure 4 (b).

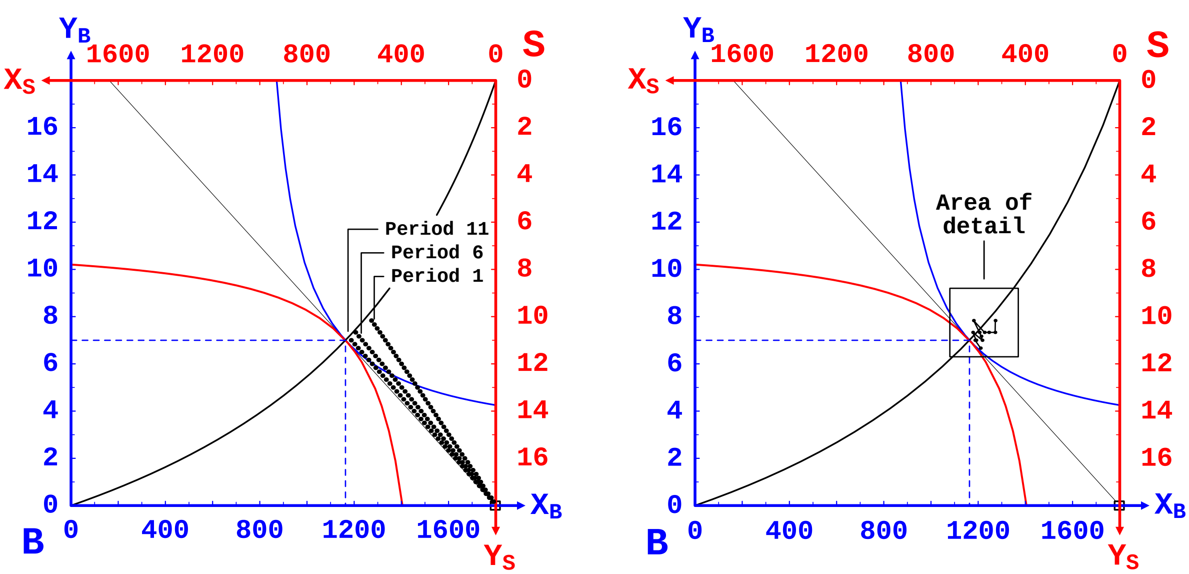

(a) Trades within periods

(b) Final allocations across periods

Figure 3: Interim allocations in three periods (left) and final allocation for all periods (right) of York1

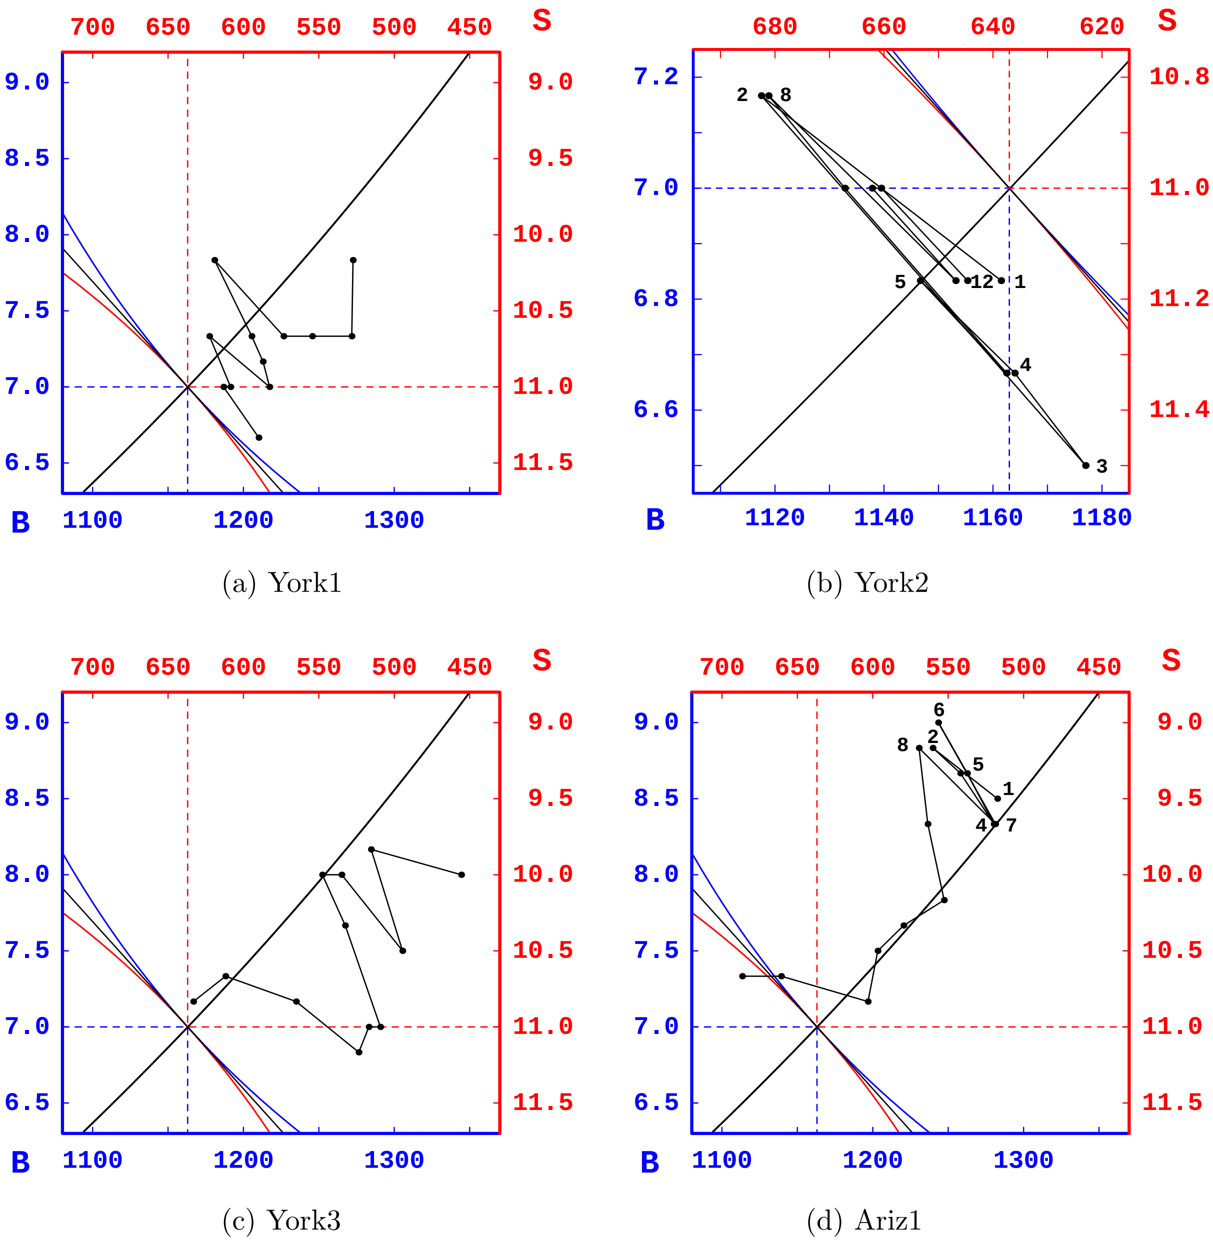

Figure 4: Edgeworth diagram details for experiment sessions.

With twelve agents and two commodities, the allocation space has twenty-four dimensions, so these diagrams substantially – though usefully – compress the data. Statistical evaluation of convergence of individual allocations to the competitive equilibrium allocations confirms the impression of convergence conveyed by these diagrams. Convergence is evaluated by measuring the distance between each trader's allocation and their competitive equilibrium allocation, and showing that these distances, averaged across traders, converge across periods.

Let $(x_{t,i}, y_{t,i})$ be the period t final allocation for agent i and let (x^*_i, y^*_i) be the equilibrium allocation for agent i. The metric used to measure the distance between the final allocation $(x_t, y_t) = ((x_{t,1}, y_{t,1}), (x_{t,2}, y_{t,2}), \ldots, (x_{t,12}, y_{t,12}))$ observed in period t and the equilibrium allocation $(x^*, y^*) = ((x^*_1, y^*_1), (x^*_2, y^*_2), \ldots, (x^*_{12}, y^*_{12}))$ is

$$d((x_t, y_t), (x^*, y^*)) = \frac{1}{12} \sum_{i=1}^{12} \frac{1}{\sqrt{2}} \left(\left(\frac{1}{91} (x_{t,i} - x^*_i) \right)^2 + (y_{t,i} - y^*_i)^2 \right)^{0.5} .$$

The rationale for this distance metric is straightforward: if trades take place at the equilibrium price $p^* = 91$ and the final allocation of each trader falls short of or exceeds the equilibrium allocation by α units, then the distance between the allocation and the equilibrium allocation is α.

With the distance $d((x_t, y_t), (x^*, y^*))$ denoted d_t, convergence is evaluated with the regression equation $d_t = d_1 \, e^{-r \ln t} \, \eta_t$ (where $\{\eta_t\}$ is a sequence of independent, identically distributed lognormal random variables). This can be expressed as the linear model $\ln d_t = \ln d_1 - r \, \ln t + \epsilon_t$.

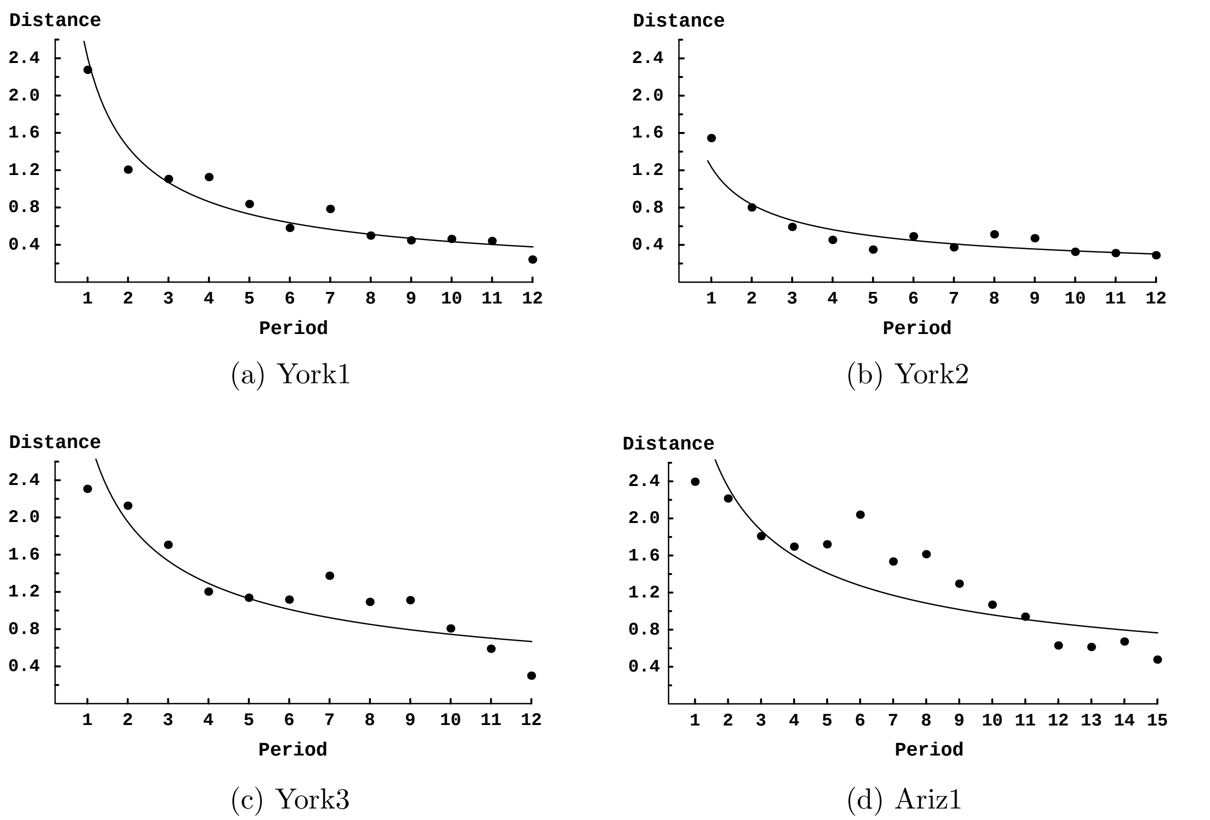

Figure 5: Convergence of allocations in four sessions.

Figure 5 shows the sequence of observed distances $\{d_t\}_{t=1}^{T}$ and also shows estimates from the regression.[6] Table 2 summarizes convergence estimates for each of the four sessions. Strong con-

[6] Although the linear model $d_t = d_1 - r\,t + \epsilon_t$ can be fit to this data, the exponential model is more appropriate, since distances cannot be negative. In addition, the average R^2 statistic for the exponential model is slightly higher than it is for a linear model.

vergence occurs in all four sessions (as measured by the estimates $\hat{r}$ and their p-values).[7] The R^2 statistics from the regression too indicate that the model captures much of the variability in the sequences $\{d_t\}_{t=1}^{T}$.

	Periods	$\hat{d}_1$	$\hat{r}$	*p-value* ($\hat{r} > 0$)	R^2	$\hat{d}_T$	$\hat{d}_T/\hat{d}_1$
York1	12	2.464	0.754	0.000	0.886	0.378	0.153
York2	12	1.235	0.566	0.000	0.826	0.302	0.245
York3	12	3.006	0.605	0.001	0.675	0.669	0.223
Ariz1	15	3.483	0.559	0.000	0.699	0.740	0.212

Table 2: This table includes estimates $\hat{d}_1$ of the initial distance from equilibrium, estimates of the rate of convergence $\hat{r}$, p-values for the hypothesis test that $r > 0$, R^2 statistics for the model, distances $\hat{d}_T$ implied by the estimates $\hat{r}$ and $\hat{d}_1$ in the last period (where $\hat{d}_T = \hat{d}_1\, e^{-\hat{r}\ln T}$), and the ratio of the implied estimate for the last period to the estimate for the first period.

4.2 Price convergence across periods: tatonnement

Across-period adjustment of average price is evaluated with the price adjustment rule from the tatonnement model $\bar{p}_{t+1} = \bar{p}_t + b\,Z(\bar{p}_t) + \epsilon_{t+1}$, where $Z(\bar{p}_t)$ is market excess demand at $\bar{p}_t$. If $\bar{p}_t = p^*$ then excess demand is zero, so the expected price in period $t+1$ is the equilibrium price, but for an adjustment rate $b > 0$ that is not too large (overshooting is possible in the discrete version of the tatonnement model), and for $\bar{p}_t \neq p^*$, the expected mean price in period $t+1$ lies between $\bar{p}_t$ and p^*.

If there is no systematic adjustment in response to excess demand ($b = 0$), this adjustment process is the random walk $\bar{p}_{t+1} = \bar{p}_t + \epsilon_{t+1}$, and if $b < 0$ the process is non-stationary, so construction of hypothesis tests requires some caution. First, in order to formulate a test of the hypothesis $b > 0$ against the null hypothesis of a random walk (if $b = 0$) or a non-stationary process (if $b < 0$), the model needs to be formulated in first differences $\Delta\bar{p}_{t+1} = \bar{p}_{t+1} - \bar{p}_t$, as

$$\Delta\bar{p}_{t+1} = b\,Z(\bar{p}_t) + \epsilon_{t+1}. \tag{4.1}$$

4.2.1 Formal adjustment test with linear approximation to excess demand

The model in equation (4.1) is similar to the "mean-reversion" model (or what in this context could naturally be called an "equilibrium reversion" model) $\Delta\bar{p}_{t+1} = \beta\,(p^* - \bar{p}_t) + \epsilon_{t+1}$. The equilibrium

[7] Nine input errors were corrected from the four sessions for the estimates reported here and in Section 6. In the four sessions there were 2291 trades, so the errors represent less than one in 254 trades. Three of the nine errors occurred when a seller omitted a digit in his ask or a buyer included an extra digit in her bid.

p^* is known to the experimenter and $\beta\,(p^* - \bar{p}_t)$ is a linear approximation to $b\,Z(\bar{p}_t)$ around p^*. Let $\bar{d}_t = \bar{p}_t - p^*$. (In Section 4.1 d_t is the distance between an allocation and the equilibrium allocation. Here $\bar{d}_t$ is the difference between the average price and the equilibrium price in period t.) Then the adjustment process can be expressed as

$$\Delta \bar{d}_{t+1} = -\beta\,\bar{d}_t + \epsilon_{t+1}. \tag{4.2}$$

The Dickey-Fuller unit root test can be applied to the model in equation (4.2) to test the null hypothesis of a random walk ($\beta = 0$) or non-stationarity ($\beta < 0$) against the alternative hypothesis that price responds in the predicted direction to deviations from equilibrium. The standard Dickey-Fuller test procedure involves inclusion of a constant term β_0, and the augmented procedure tests lagged dependent variables. The augmented model is tested, and insignificant lags are eliminated. For the mean price sequences from the four experiment sessions, no lagged price change is significant, so the model $\Delta d_{t+1} = \beta_0 - \beta\,\bar{d}_t + \epsilon_{t+1}$ is estimated. The t-statistics for $\hat{\beta}_0$ fall within the range $(-2.179, 2.179)$ for sessions York1, York2, and York3, so the hypothesis of a constant term in the adjustment process is rejected for these three sessions using a standard t-test.[8] The null hypothesis of no adjustment due to price deviation from equilibrium is then tested in the model of equation (4.2), which doesn't include a constant term. For that model, the critical value of the t-statistic is the Dickey-Fuller test statistic $\boldsymbol{\tau} = -1.95$ for a test at the 95% confidence level. The second row for each session in table 3 shows estimates for this model and the t-statistic for the test of the null hypothesis that $\beta = 0$. For sessions York1, York2, and York3, the null hypothesis is rejected: adjustment in each case is significant and in the direction of equilibrium.[9] Results from session Ariz1 are somewhat mixed. The constant term is significant, and in the model with the constant term included, it is not possible to reject the null hypothesis that prices follow a random walk or a non-stationary process, since the adjustment term has the wrong sign. If attention is restricted to the last eight periods in Ariz1, when adjustment took place, the result differs. In periods $8 - 15$, the constant term remains significant, but in the model without a constant term, adjustment is significant in the model. Elliott, Rothenberg, and Stock [1996] demonstrate that Dickey-Fuller tests have very low power when the product of the sample size and the adjustment

[8] The average price in the first period of York2 was $\bar{p}_1 = 93.44$, which is very near the competitive equilibrium price $p^* = 91$. The average price moved away from the competitive equilibrium price in period 2 to $\bar{p}_2 = 95.23$, moved further away in period 3 to $\bar{p}_3 = 95.85$, and then began to slowly adjust back to the equilibrium price. Although the average price movements in these two periods were small, they moved in the opposite direction from the predicted adjustment in response to excess demand. It is remarkable that price adjusted between periods 4 and 12 toward the competitive equilibrium: excess demand at the mean price in period 3 was $Z^Y(95.85) = -0.799$, or -0.133 on a per capita basis. With periods 1 and 2 included, it is not possible to reject the hypothesis of a random walk.

[9] The general Dickey-Fuller test procedure is described in Enders [1995]. See in particular equations (4.9), (4.10), figure 4.7. Enders provides critical values for the Dickey-Fuller test statistics in his Table A.

parameter is small.[10] In this situation, the null hypothesis is frequently accepted when it is in fact false. Given this situation, it is noteworthy that the null hypothesis can be rejected for the first three sessions, and for the last eight periods of the session Ariz1.

	Periods	β_0	$-\beta$	t-stat	Test
York1	1 – 12	−0.038	−0.175	−0.086	$\beta_0 = 0$
			−0.172	−11.101	$\beta = 0$
York2	3 – 12	0.232	−0.096	0.366	$\beta_0 = 0$
			−0.041	−2.432	$\beta = 0$
York3	1 – 12	3.041	0.009	1.814	$\beta_0 = 0$
			−0.123	−4.334	$\beta = 0$
Ariz1	1 – 15	5.973	0.166	5.306	$\beta_0 = 0$
			−0.066	−2.046	$\beta = 0$
Ariz1	8 – 15	4.873	0.047	0.205	$\beta_0 = 0$
			−0.198	−3.246	$\beta = 0$

Table 3: Regression coefficients and statistics for the linear adjustment rule.

4.2.2 Informal adjustment test with excess demand

Equation (4.2) conforms exactly to the model specification in the Dickey-Fuller tests for a unit root, whereas equation (4.1) conforms to the microeconomic model of price adjustment. Equation (4.1) though can be put into a form similar to equation (4.2). First, use the change of variable $\bar{d}_t = \bar{p}_t - p^*$ to write equation (4.1) as

$$\Delta \bar{d}_{t+1} = b\, Z(\bar{d}_t + p^*) + \epsilon_{t+1}. \tag{4.3}$$

In equation (4.2) the regressor $\bar{d}_t$ passes through the origin with slope one. In equation (4.3) the regressor $Z(\bar{d}_t + p^*)$ passes through the origin but its slope there is -0.1713. The function $\tilde{Z}(\bar{d}_t) = -Z(\bar{d}_t + p^*)/0.1713$ passes through the origin with slope one, so

$$\Delta \bar{d}_{t+1} = -b'\, \tilde{Z}(\bar{d}_t) + \epsilon_{t+1} \tag{4.4}$$

is equivalent to equation (4.1) and it is tangent to equation (4.2) at $\bar{d}_t = 0$ (if $b' = \beta$).

The Dickey-Fuller test statistic is not formally applicable to equation (4.4), but the regressor $\tilde{Z}(\bar{d}_t)$ approximates the regressor $\bar{d}_t$ in equation (4.1) from the formal statistical model, so the test can be carried out with equation (4.4), albeit informally.

[10] A central argument in this paper is that within-period adjustment exceeds across-period adjustment. The power of Dickey-Fuller tests is much higher in tests of the Hahn process than in tests of tatonnement adjustment, since the sample size for the Hahn process approximately equals the number of trades rather than the number of periods. This provides an econometric rationale that complements other arguments in the paper for within-period analysis.

	Periods	b_0	$-b'$	t-stat	Test
York1	1 - 12	0.352	-0.124	1.009	$b_0 = 0$
			-0.143	-11.175	$b' = 0$
York2	3 - 12	0.206	-0.093	0.336	$b_0 = 0$
			-0.043	-2.431	$b' = 0$
York3	1 - 12	2.677	0.007	2.088	$b_0 = 0$
			-0.089	-4.064	$b' = 0$
Ariz1	1 - 15	5.463	0.105	6.018	$b_0 = 0$
			-0.041	-1.702	$b' = 0$
Ariz1	8 - 15	4.780	0.035	7.081	$b_0 = 0$
			-0.141	-2.548	$b' = 0$

Table 4: Regression coefficients and statistics for the non-linear adjustment rule.

Table 4 reports estimates and t-statistics for the model in equation (4.4) in which price adjusts in response to excess demand, rather than in response to deviations from equilibrium. All conclusions remain the same with the alternative formulation. In this formulation too, the lack of adjustment in the first seven periods of Ariz1 seems to be the reason that the null hypothesis $\beta \leq 0$ cannot be rejected. (This lack of adjustment is most apparent in figure 8.) When the test is performed on periods 8 – 15, the adjustment term is, at $b' = 0.141$, as strong as in any session, and its t-statistic of -2.548 is well below the Dickey-Fuller test statistic $\boldsymbol{\tau} = -1.95$. Since lagged price changes are rejected in the data, the lack of adjustment between periods 2 and 7 does not result from serial correlation.[11] It must be the result of a structural change. In fact, evidence of that structural change is apparent in the data: beginning near the end of period 6, one buyer recognized that she could increase her utility by raising her bid above the prevailing price. Figure 1 shows data from the end of period 6, including the market queue. All prices in the period were at 60, 61, or 62 until Buyer 2 paid 65 near the end of the period and then paid 70 in the final trade of the period. Near the end of the period, she raised her final bid to 79. At the end of period 7 three other buyers imitated her behavior. Starting in period 8, the structural change was fully in effect. (See figure 8.) It is interesting and notable that the econometric technique rejects serial correlation in errors as the explanation and points instead (via comparison between the test statistics for the full sample and the subsample of periods) to a structural change, which examination of individual behavior supports.

[11] In these models, the presence of lagged price changes in the adjustment process and the presence of serial correlation in the error structure are indistinguishable. Appendix B demonstrates that a model with first-order serial correlation is equivalent to a model with a single lagged price change in the adjustment function. In general, a model with $AR(n)$ errors is equivalent to a model with n lagged price adjustment terms. (See Davidson and McKinnon [1993, pp. 710 – 711].)

4.3 Comparison of across-period and within-period price adjustment

Section 4.1 demonstrates that allocations converge to the competitive equilibrium allocation across trading periods. Section 4.2 demonstrates that average prices converge to the competitive equilibrium price with tatonnement price adjustment. This section compares the extent of across-period price adjustment to the extent of within-period price adjustment and demonstrates that across-period price changes are typically smaller than within-period price changes, which motivates the analysis of within-period price adjustment in Sections 5 and 6.

For periods $t = 2, 3, \ldots, T$, the ratio $A_t^a = 1 - |\bar{p}_t - p^*|/|\bar{p}_{t-1} - p^*|$ measures price adjustment across periods. Denote the number of trades in period t by K_t. Within-period price adjustment for periods $t = 1, 2, \ldots, T$ is measured as the ratio $A_t^w = 1 - \left|\sum_{k=K_t-4}^{K_t} (p_{t,k} - p^*)\right| \Big/ \left|\sum_{k=1}^{5} (p_{t,k} - p^*)\right|$.

Three cases help interpret these measures of across-period and within-period price adjustment. The first case involves only within-period adjustment, the second considers equal within-period and across-period adjustment, and the final one involves only across-period adjustment.

(1) If price adjusts within each period so that $\sum_{k=K_t-4}^{K_t} (p_{t,k} - p^*) = (1-\alpha) \sum_{k=1}^{5} (p_{t,k} - p^*)$, and if there is no adjustment across periods, then $A_t^a = 0$ and $A_t^w = \alpha$.

(2) If price adjusts geometrically within each period so that $(p_{t,K_t} - p^*) = (1-\alpha)(p_{t,1} - p^*)$, and if the first price in period $t+1$ is equal to the last price in period t, then $A_t^a = \alpha$ and $A_t^w \doteq \alpha$.

(3) If price in each period is constant and $(p_{t,k} - p^*) = (1-\alpha)(p_{t-1,k} - p^*)$, then $A_t^a = \alpha$ and $A_t^w = 0$.

Average across-period price adjustment A_t^a for the four sessions was $\bar{A}^a = 0.078$ and average within-period adjustment A_t^w was $\bar{A}^w = 0.238$.[12] Although the tatonnement model tracks across-period price changes, these average price changes are most likely a response to large within-period price changes. The next section explores the predictions of the Hahn process model of within-period price adjustment in order to better understand price adjustment.

5 Hahn process dynamics

In the four experiment sessions, prices and allocations converge across periods to the competitive equilibrium. Since the magnitude of within-period price changes exceeds that of across-period price

[12] In the first two periods of York2, the mean price moved away from the equilibrium price. In the fifteenth period of Ariz1, the mean price overshot the equilibrium price slightly, going from $\bar{p}_{14} = 90.1$ to $\bar{p}_{15} = 93.6$. Both these movements decrease the measures of within-period and the across-period price changes. Without these three periods, the measures are $\bar{A}^a = 0.146$ for across-period price adjustment and $\bar{A}^w = 0.347$ for within-period price adjustment. With or without these three outliers, these measures indicate that there is substantially more price adjustment within periods than across periods.

changes, within-period price dynamics are evaluated with the Hahn process model of disequilibrium price adjustment.

This section begins, in Section 5.1, with a brief description of the continuous version of the Hahn process formulated in Hahn and Negishi [1962]. Section 5.2 adapts the model to the context of the market experiment with trade in discrete units and discrete price adjustment. The experiment includes several replications of the exchange economy, so Section 5.3 augments the Hahn process model with a simple rule for across-period price adjustment.

5.1 Hahn process description and notation

In an exchange economy with m commodities and n agents, each with a utility function $u_i(x, y)$ and an endowment ω_i, prices in the Hahn process adjust in response to excess demand. Unlike tatonnement though, in the Hahn process, trade takes place as soon as a price is announced. In the continuous version of the Hahn process, the initial price vector is held constant during some time interval $[0, h]$, and subsequently (for $t > h$) prices adjusts continually in response to excess demand according to the rule $\dot{p}(t) = c\, Z^{Y}(p(t)|x(t))$, where $x(t)$ is the allocation that results from all trades that have occurred along the price path $(p(s))_{s\in(0,t)}$. Trades of the goods follow some rule $\dot{x}_i(t)$ that satisfies three conditions: (1) the budget of each agent i balances $(p(t) \cdot \dot{x}_i(t) = 0)$; (2) the net flow of each commodity balances $(\sum_{i=1}^{n} \dot{x}_i(t) = 0)$; and (3) the executed trade increases the utility of at least one agent and does not decrease the utility of any agent $(u_i(x_i(t) + \dot{x}_i(t)) \geq u_i(x_i(t))$ for all i with strict inequality for at least one agent). Hahn and Negishi [1962] and Arrow and Hahn [1971, Chapter 13] demonstrate that this price adjustment process converges to a Pareto optimal allocation for any adjustment rate c and any trade process that satisfies the three indicated conditions.[13]

5.2 Discrete Hahn process price adjustment

A discrete version of this process adapts the model to the experiment. The discrete version with two commodities operates simply. At the announced price $p_{t,k}$ some buyer with excess demand at that price transfers $p_{t,k}$ units of the numeraire commodity to a seller, who in return transfers one

[13] Their main claim though is that as adjustments approach their continuous limit, prices and allocations converge to a competitive equilibrium of the original economy. In the Scarf example, if price adjusts slowly and the trade quantity at each announced price is relatively large, the limit of both the price and the allocation differ from the competitive equilibrium allocation. On the other hand, if price adjustment is rapid and trade quantities are small, the process can converge, even in a single trading period, to the competitive equilibrium of the Scarf example. Nevertheless, even if the theorem were stated and proven correctly, the result would only be applicable to economies much larger than the one in the experiment, due to the assumption in the Hahn process model that trades are very small relative to equilibrium net trades. In the experiment, convergence is to a Pareto optimal allocation within periods; convergence to the competitive equilibrium price and allocation takes place across periods.

unit of the other commodity to the buyer. With this adjustment process, rules (1) and (2) are met by construction; rule (3) is imposed in simulations by ending trade when there is no Pareto improving trade at the announced price.[14]

The excess demand of each agent depends on his current allocation and on the market price, so implementation of the model requires an expression for the allocation of each trader in terms of his initial endowment and the interim allocation that results after each trade. Suppose that from the sequence $P_{t,k} = (p_{t,1}, p_{t,2}, p_{t,3}, \ldots, p_{t,k})$ of the first k trades in period t, agent i has taken part in k_i trades. These k_i trades by agent i can be denoted $P^{(i)}_{t,k} = (p_{t,i_1}, p_{t,i_2}, \ldots, p_{t,i_{k_i}})$. The allocation for buyer $i \in B$ can be written in terms of $P^{(i)}_{t,k}$ as $\left(x_{t,k}(P^{(i)}_{t,k}),\, y_{t,k}(P^{(i)}_{t,k})\right) = \left(1800 - \sum_{j=1}^{k_i} p_{t,i_j},\, k_i\right)$. The allocation for seller $i \in S$ is $\left(x_{t,k}(P^{(i)}_{t,k}),\, y_{t,k}(P^{(i)}_{t,k})\right) = \left(\sum_{j=1}^{k_i} p_{t,i_j},\, 18 - k_i\right)$.

In the discrete version of the Hahn process, price adjusts after each trade in response to aggregate excess demand. After trade k in period t, price adjusts so that

$$p_{t,k+1} = p_{t,k} + c \sum_{i=1}^{12} Z_i^Y \left(p_{t,k} \,\Big|\, \left(x_{t,k}(P^{(i)}_{t,k}),\, y_{t,k}(P^{(i)}_{t,k})\right)\right). \tag{5.1}$$

The estimation procedure in Section 6.2 uses the expression for excess demand from the right side of equation (5.1) as a regressor. Simulation uses estimates from the augmented version of this model, which includes statistically significant lagged price changes.

5.3 Initial prices in periods t = 2, 3, . . . , T

This section examines the distribution of initial prices in period $t+1$ in terms of the range of prices in period t. The Hahn process estimation procedure does not capture this aspect of the adjustment process, so it is examined separately and combined with estimates of within-period adjustment in a twelve-period simulation in Section 6.3.

The within-period price path predicted by the Hahn process has two free parameters: the initial price $p_{t,1}$ and the adjustment rate c. The model is adapted in this paper to account for the replications of the economy across trading periods by including a simple rule for setting the initial price in periods $t = 2, 3, \ldots, T$. The initial price in period t is assumed to be a weighted average of the maximum price and the minimum price from period $t-1$:

$$p_{t,1} \equiv (1-a)\, \min\{p_{t-1,1}, p_{t-1,2}, \ldots, p_{t-1,K_{t-1}}\} + a\, \max\{p_{t-1,1}, p_{t-1,2}, \ldots, p_{t-1,K_{t-1}}\}. \tag{5.2}$$

Forty-seven observations of a are obtained from the fifty-one periods in four experiment sessions. In the four sessions, the median value of a was $a_{24} = 0.477$. The mean was $\hat{a} = 0.476$, its sample

[14] This rule may not reach a Pareto optimal allocation because it can stop at the net demand of one agent. Other ending rules that are consistent with condition (3) are also possible. A rule that repeatedly changes the price if no Pareto improving trade is possible at the current price will eventually reach a Pareto optimal allocation.

variance was $s^2 = 0.083$, and its minimum and maximum values were zero and one. These statistics are almost indistinguishable from a uniform distribution on $[0, 1]$.

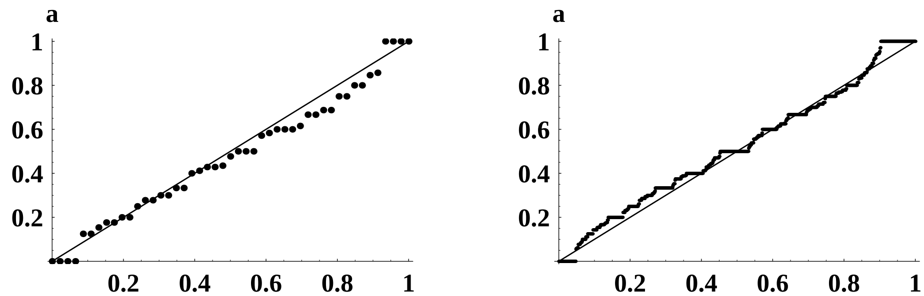

Figure 6: Distribution of initial prices $p_{t,1}$ in the four sessions (left) and simulation (right).

The distribution of a is shown on the left side in figure 6. The data on the right side of the figure are simulated values of a for the forty-seven individual period price ranges from the experiment sessions. The data were generated by taking each of these forty-seven price ranges, selecting an (integer-valued) initial price for the next period uniformly from that interval, and then calculating the statistic a for that initial price. This was replicated twenty times to get a simulation of the distribution of initial prices. Based on the distribution of a from the experiment and the simulated distribution, it seems reasonable to augment the Hahn process price adjustment model for replicated market environments with the initial price in period $t+1$ selected randomly from a uniform distribution on the interval between the minimum and the maximum price in period t.[15]

6 Hahn process estimation and simulation

With error terms added to predicted price adjustments, prices in period t evolve according to the equation $p_{t,k+1} = p_{t,k} + c\,Z^{Y}(p_{t,k}) + \epsilon_{t,k+1}$. As with the tatonnement model, if $c = 0$ the adjustment process follows a random walk and if $c < 0$, the process is non-stationary. Consequently, the test for $c > 0$ must be conducted in first differences $\Delta p_{t,k+1} = p_{t,k+1} - p_{t,k}$ as

$$\Delta p_{t,k+1} = c\,Z^{Y}(p_{t,k}) + \epsilon_{t,k+1}. \tag{6.1}$$

[15] Easley and Ledyard [1993] construct their model of double auction bidding behavior from their observation – based on their examination of data from double auction trading experiments – that in market environments with induced costs and values prices in period $t+1$ typically fall within the observed price range from period t. My claim is similar, but has two differences. I make the weaker claim that only the initial price in period $t+1$ falls in the range of period t prices, and I make the more specific assumption that it is drawn uniformly from that interval. The two approaches are similar: both are based on observations from double auction experiments. The different economic environment that I consider leads to my restriction of the assumption to initial prices, and I make the more specific distributional assumption based on statistical properties of initial prices.

6.1 Hahn process estimation with linear approximation to excess demand

As with the test of the tatonnement model, the model in equation (6.1) does not fit exactly into the framework of the unit-root tests. With the linear approximation $\gamma\,(p^* - p_{t,k})$ to $c\,Z^Y(p_{t,k})$ in equation (6.1), the adjustment equation can be written $\Delta p_{t,k+1} = \gamma\,(p^* - p_{t,k}) + \epsilon_{t,k+1}$. Since it is deviations from equilibrium that lead to price adjustment, $d_{t,k} = p_{t,k} - p^*$ can be substituted into this equation to get the equivalent adjustment model $\Delta d_{t,k+1} = -\gamma\, d_{t,k} + \epsilon_{t,k+1}$. This is the version of the Dickey-Fuller test without a constant or a trend term. In the general test procedure, we can add a constant term and lagged price changes. The adjustment model doesn't include a constant term, so we expect that the constant term will be rejected. The model also does not include lagged price changes, though that seems plausible.[16] With a constant term and lagged price adjustment, the model is

$$\Delta d_{t,k+1} = \gamma_0 - \gamma\, d_{t,k} + \sum_{j=1}^{q} \beta_j\, \Delta d_{t,k+1-j} + \epsilon_{t,k+1}. \tag{6.2}$$

In this form, the hypothesis test $\gamma > 0$ is carried out by comparing the t-statistic from the parameter estimate to the critical value of the Dickey-Fuller test statistic for a unit-root process.[17]

The number of lags can be examined either by a general-to-specific sequential rule or the Bayesian information criterion (BIC). The general-to-specific sequential rule begins with a large number of lags (set at $q = 8$ here), and then insignificant lags are eliminated with a standard t-test. For the model with a constant and approximately 500 sample points, if the absolute value of the t-statistic is greater than 1.96 the hypothesis that a lagged price adjustment parameter β_j equals zero can be rejected at the 5% significance level, using a standard t-test.

This test was used to eliminate insignificant lags. [18] (Only significant lags are reported in tables 5 and 6.) Once insignificant lags are eliminated, the hypothesis that the constant term is zero is tested with a standard t-test. Since the constant term is significant in York2, the test statistic $\boldsymbol{\tau_\mu} = -2.87$ is used to reject the hypothesis of no adjustment.[19]

[16] As outlined in footnote 10 and detailed Appendix B, lagged price changes in the adjustment process and serial correlation in the error structure are indistinguishable.

[17] The procedure for these tests is similar to the tests with the tatonnement model in Section 4.2, but lags are significant in the within-period adjustment process. For these tests, see equations (4.12) and (4.13) and figure 4.7 in Enders [1995].

[18] In general, the BIC is consistent while the general-to-specific rule has positive probability of over-fitting. Even so, the results for the four exchange economy sessions are identical with the two procedures.

[19] In York2 excess demand was small, and adjustment was in the right direction, but the impetus to adjustment was so subtle that it is unsurprising that the constant term is significant (i.e., prices converged statistically to a price slightly different from the equilibrium price). If the first period is not included, and the equilibrium price is taken as $p_e = 94$ rather than $p^* = 91$, then the constant term is not significant while the adjustment term remains significant.

In the three sessions for which the constant term is rejected, the model is estimated without a constant to test for significance of the adjustment term. In each of these sessions, the adjustment term is significant with a p-value below 5% (Dickey-Fuller test statistic $\boldsymbol{\tau} = -1.95$). In fact, in sessions York1 and York3, it is significant with a p-value below 1% (Dickey-Fuller test statistic $\boldsymbol{\tau} = -2.58$). These results are all summarized in Table 5, which shows parameter estimates for the model in equation (6.2). In session York2, which has a significant constant term, the null hypothesis of a unit root (no adjustment) is rejected. (The critical value of the Dickey-Fuller test with a constant term is $\boldsymbol{\tau_\mu} = -3.44$) at a p-value of 1%.)

	N	γ_0	$-\gamma$	β_1	β_2	β_3	β_4	β_5	t-stat	Test
York1	511	0.171	−0.034	−0.692	−0.601	−0.463	−0.300	−0.287	0.83	$\gamma_0 = 0$
	—		−0.045	−0.680	−0.590	−0.453	−0.293	−0.283	−5.45	$\gamma = 0$
York2	483	0.721	−0.198	−0.244					5.56	$\gamma_0 = 0$
	—	0.721	−0.198	−0.244					−6.41	$\gamma = 0$
York3	527	0.420	0.005	−0.721	−0.443	−0.437	−0.219		1.78	$\gamma_0 = 0$
	—		−0.013	−0.699	−0.422	−0.419	−0.206		−2.90	$\gamma = 0$
Ariz1	718	−0.293	−0.022	−0.693					−1.13	$\gamma_0 = 0$
	—		−0.011	−0.699					−2.13	$\gamma = 0$

Table 5: Parameter estimates and t-statistics for the linear adjustment rule in equation (6.2).

6.2 Hahn process estimation with excess demand

Equation (6.2) has the merit that it is specified so that the Dickey-Fuller tests apply to it. The model too is very similar to the Hahn process adjustment model, since $Z^Y(p_{t,k})$ is nearly linear in a fairly large region around the equilibrium price p^*. As a consequence of this, the results from the model in equation (6.2) should be similar to the results from the model

$$\Delta p_{t,k+1} = a_0 + c\, Z^Y(p_{t,k}) + \sum_{j=1}^{q} b_j\, \Delta p_{t,k+1-j} + \epsilon_{t,k+1}. \tag{6.3}$$

Comparison is facilitated if, as with the unit root test for the tatonnement adjustment model, we make the change of variable $d_{t,k} = p_{t,k} - p^*$ and then set $\tilde{Z}^Y(d_{t,k}) = -Z^Y(d_{t,k})/0.1713$. Then the regressor $\tilde{Z}^Y(d_{t,k})$ in

$$\Delta d_{t,k+1} = a_0 - c'\, \tilde{Z}^Y(d_{t,k}) + \sum_{j=1}^{q} b_j\, \Delta d_{t,k+1-j} + \epsilon_{t,k+1} \tag{6.4}$$

has the same slope at $d_{t,k} = 0$ as the regressor in equation (6.2), so that estimates from the two are comparable. In fact, all of the results for this model are identical to those for the model in equation (6.2). All of the results in table 6 are interpreted similarly to those in table 5.

	N	a_0	$-c'$	b_1	b_2	b_3	b_4	b_5	t-stat	Test
York1	511	0.199	−0.028	−0.678	−0.589	−0.455	−0.295	−0.284	1.12	$a_0 = 0$
	—		−0.037	−0.660	−0.571	−0.439	−0.282	−0.278	−5.54	$c' = 0$
York2	483	0.535	−0.165	−0.245					4.90	$a_0 = 0$
	—	0.535	−0.165	−0.245					−5.95	$c' = 0$
York3	527	0.233	0.003	−0.707	−0.431	−0.428	−0.212		1.19	$a_0 = 0$
	—		−0.010	−0.686	−0.412	−0.412	−0.202		−3.21	$c' = 0$
Ariz1	719	−0.186	−0.012	−0.687					−0.80	$a_0 = 0$
	—		−0.007	−0.694					−2.09	$c' = 0$

Table 6: Parameter estimates and t-statistics for model in equation (6.4)

6.3 Construction of predicted price sequences

The sequence of predicted price changes $\Delta p_{t,k+1}$ for $k = 1, 2, 3, \ldots, K_t$ from equation (6.4) with parameter estimates substituted into the equation can be used to generate the predicted price path from each session. The predicted price sequence is $p_{t,1}$, $p_{t,2} + \Delta p_{t,2}$, $p_{t,3} + \Delta p_{t,3}$, $\ldots$, $p_{t,K_t} + \Delta p_{t,K_t}$. Figures 7 and 8 show the trade prices and the estimated adjustment function in each session. In each graph, successive periods are separated by a vertical line, trade prices are shown as dots, and the number of trades is indicated for each period. The equilibrium price is shown as a solid horizontal line at $p^* = 91$. The estimated adjustment path is shown as a solid path in each graph (though the estimated path is often difficult to distinguish from the observed price path). Prices demonstrate a remarkably tendency toward the equilibrium, even in York2, which was so close to the equilibrium in the third period that the excess demand on a per capita basis was less than -0.133 units.

6.4 Hahn process simulation

Hahn process simulation with discrete trades in two commodities and replication of trading periods is fully specified by (1) an initial price $p_{1,1}$ in period 1; (2) the price adjustment rule in equation (6.3); and (3) the initial price in periods $t = 2, 3, \ldots, T$ in equation (5.2). The price path through T periods is only a function of the initial price $p_{1,1}$ in period 1, the adjustment rate c, the coefficients β_j on lagged price changes, and the random terms $\epsilon_{t,k}$ and across-period adjustment terms a.

For a fixed value of c, a predicted price path for period t is determined starting from the initial price $p_{t,1}$. The predicted second price $p_{t,2}$ is determined based on $p_{t,1}$ and the adjustment rule in equation (6.3). Subsequent predicted prices are determined iteratively.

The Hahn process model and its adaptation to the replicated economies of experiments can be brought together in a simulation of the process. Figure 9 shows the path of prices within and

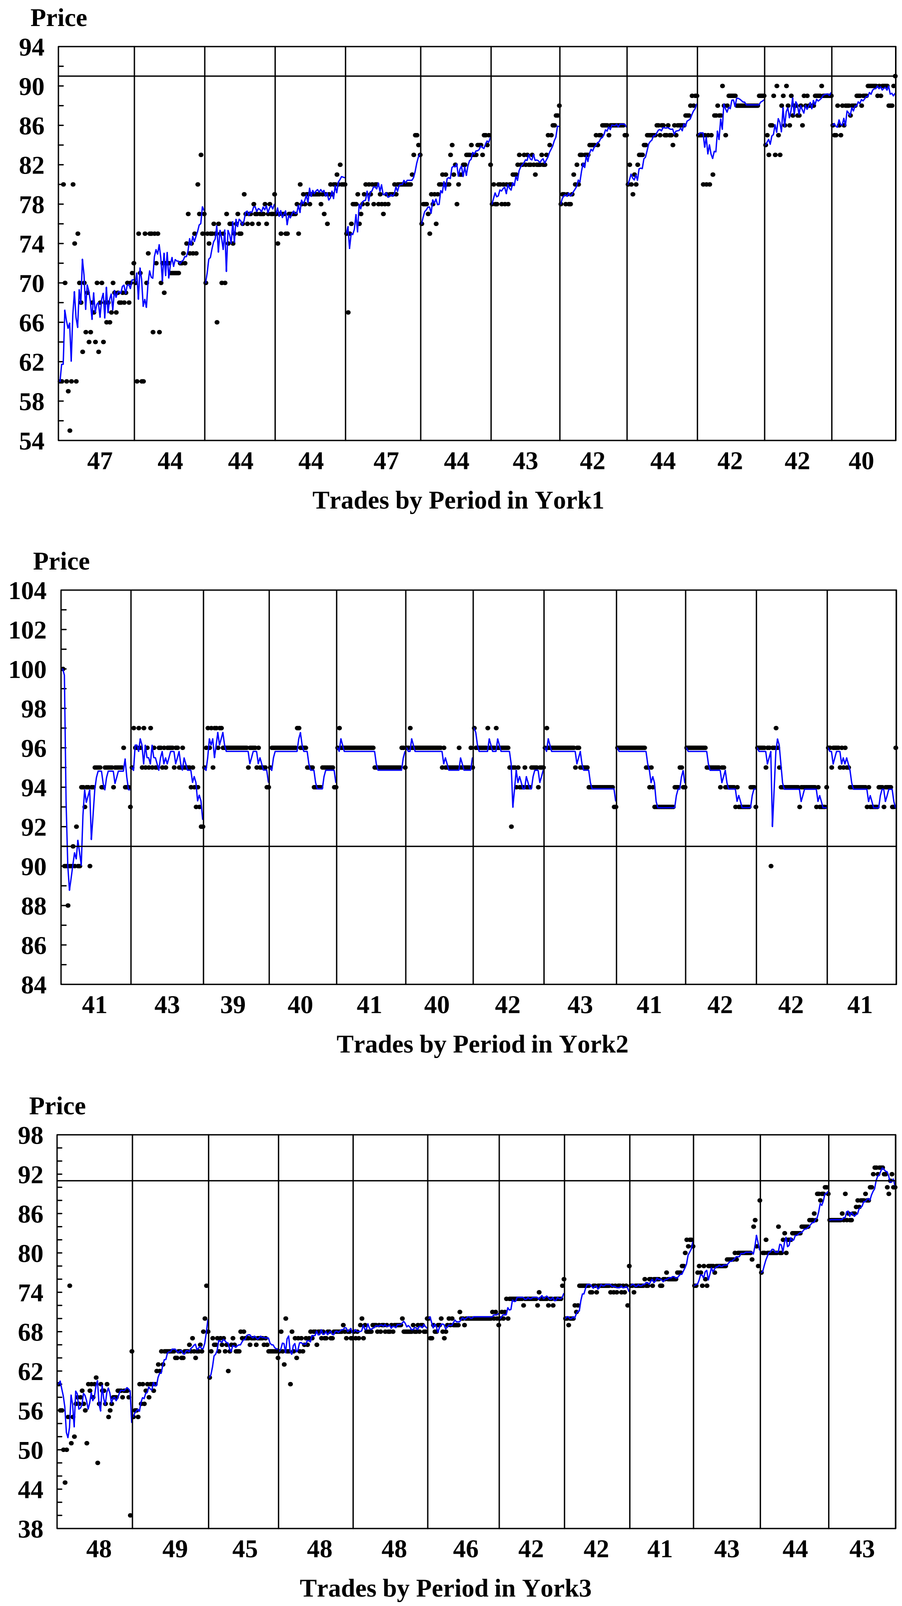

Figure 7: Prices and estimated adjustment in York1, York2, and York3.

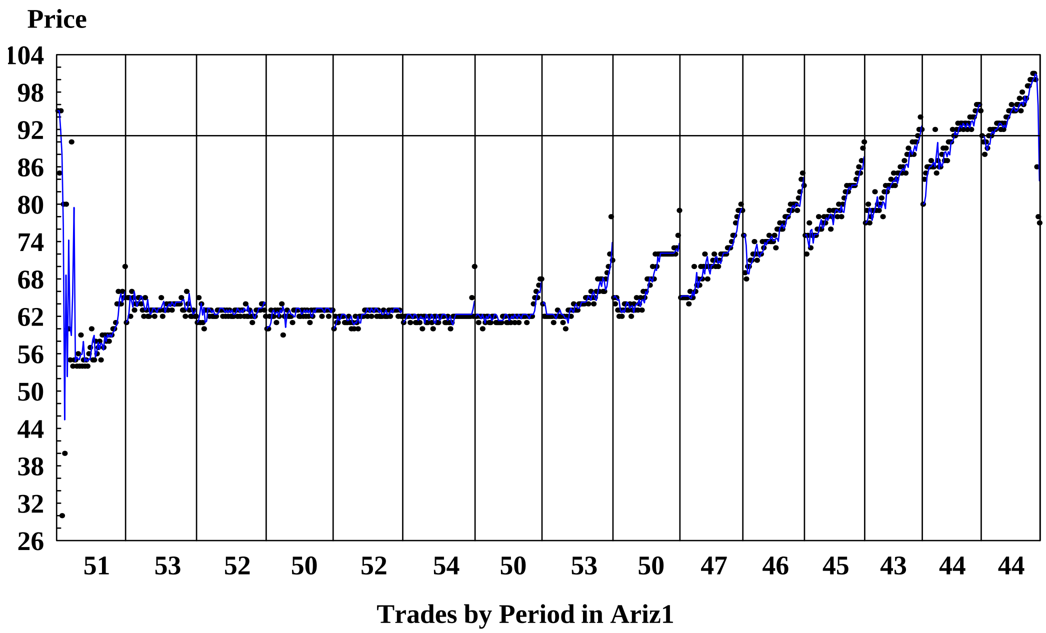

Figure 8: Prices and estimated adjustment in Ariz1.

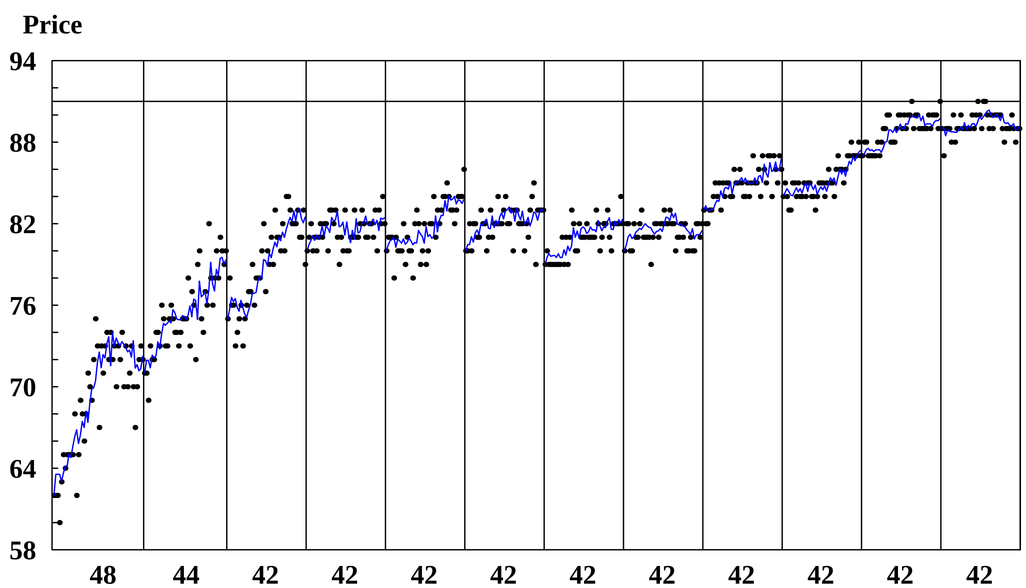

Figure 9: Simulation of price path within and across periods.

across trading periods from a simulation of this model. In this simulation, within-period prices are determined by equation (6.3), with parameter estimates from session York1. The initial price in period $t+1$ is selected randomly from the interval of period t prices, and the standard deviation of the errors decreased geometrically across periods according to the equation $\sigma_t = 2 \cdot 0.9^{t-1}$. Comparison of this figure with figure 7 (a) suggests that the Hahn process price adjustment model captures many key features of the price path.

7 Conclusions

The pure exchange economy is the fundamental model in the neoclassical theory of value. Empirical assessment of the Generalized Axioms of Revealed Preference (GARP), as in Cox [1997], demonstrates that choices over a small number of commodities are consistent with the existence of a regular preference ordering or utility function. The experiments reported in this paper demonstrate that when agents with (induced) regular utility functions trade in an exchange economy, prices and allocations converge across trading periods to its competitive equilibrium. Consequently, the two key assumptions in the neoclassical theory of value – agents have consistent preference orderings and exchange by agents with consistent preference orderings (or utility) leads to a competitive equilibrium allocation – are supported by experimental evidence, at least when these two assumptions are evaluated separately.

Equilibrium predictions though are not the whole story. Price dynamics in an exchange economy have a complex structure. Exchange economy experiments exhibit both across-period and within-period convergence. Comparison of price adjustment across periods and price adjustment within periods demonstrates that most price adjustment occurs within each trading period. Consequently, disequilibrium models that predict price adjustment within trading periods are essential in order to understand the dynamics of convergence in general equilibrium models. Previous experiments with induced values and costs did not reveal the full complexity of the convergence process.

The augmented Hahn process price adjustment model estimated in this paper is broadly consistent with price adjustment in partial equilibrium environments that have been studied in numerous CDA experiments beginning with Smith [1962]. Experimentally, a partial equilibrium market environment is created in the laboratory by inducing values for buyers and costs for sellers. Gjerstad and Shachat [2007] demonstrate that any partial equilibrium environment can be generated from an exchange environment with quasi-linear utility functions, so the price adjustment model estimated in this paper should apply also to adjustment in CDA experiments with induced values and costs, since these too can be viewed as exchange economy environments. In partial equilibrium market experiments, during early periods prices typically come from one side of the equilibrium in a wide arc that approaches the equilibrium price. In each successive period prices tend to follow a flatter arc, with the initial price closer to the equilibrium price than in the previous period, and the final price in each period approaches the equilibrium price. In these partial equilibrium market environments, even when units trade at prices away from the competitive equilibrium price, the equilibrium price of the economy that remains after a disequilibrium trade is typically the same as or very close to the equilibrium price of the original economy. As a result, trade prices converge to the equilibrium price at the end of the first period. In the second period, prices typically fall in the range of period 1 prices (as in the Easley-Ledyard model). As this process proceeds across periods, prices converge to the equilibrium price, in the strict sense that prices are near the equilibrium price throughout the period, not just at the end of the period.[20]

The exchange economy experiment provides strong evidence for convergence to competitive equilibrium, and also demonstrates the importance of within-period price adjustment. A third issue also motivates the experiment. Experimental evidence on price adjustment in standard partial equilibrium market settings led eventually to development of models of price adjustment based on the incentives of the individual agents. Models of this sort include Wilson [1987], Friedman [1991], Easley and Ledyard [1993], Gjerstad and Dickhaut [1998], and Gjerstad [2007]. The objective of this line of research has been to gain insight into the processes that lead markets to their equilibrium prices and allocations by examining the interactions among sellers and buyers. Extension of this line of research to the general equilibrium context requires a base of facts to be explained. The experiment reported in this paper examines the most elemental general equilibrium model, and as such is a natural starting point for models of price formation in general equilibrium markets.

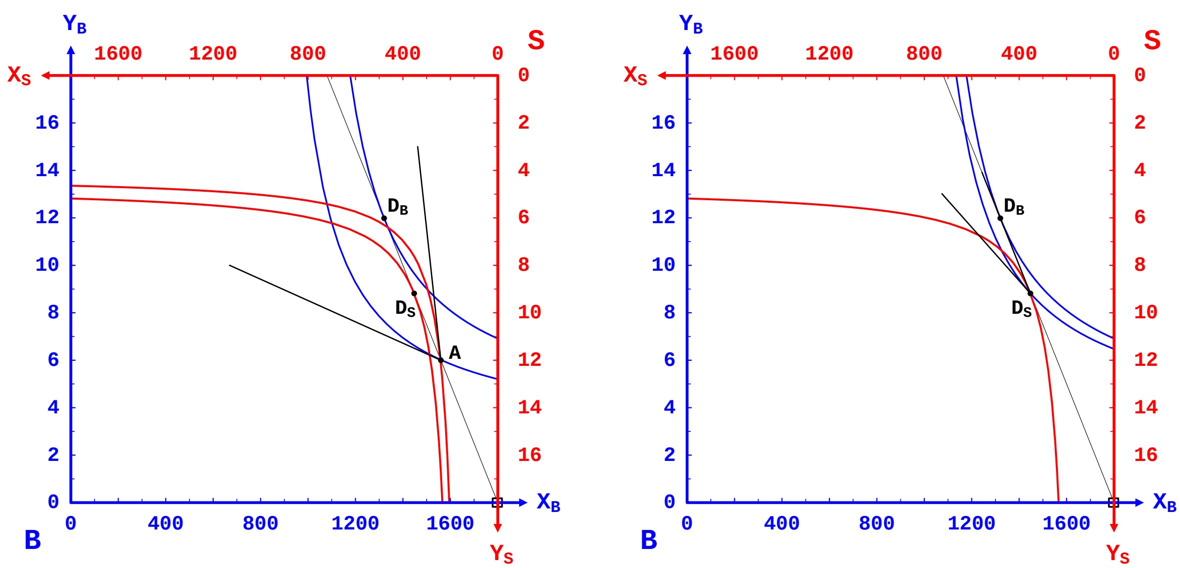

(a) Near middle of trading period (b) Near end of trading period

Figure 10: Cone of Pareto improving prices and individual demands

In general, individual incentives vary much more than the global conditions summarized in excess demand. A simple example illustrates this point. If prices are constant, then excess demand does not change as trades are conducted. In figure 10 (a), suppose that the current allocation

[20] This issue is discussed in Gjerstad and Shachat [2007].

is at point A. Then the demand of agent B is D_B and the demand of agent S is D_S. Excess demand for Y is the vertical component of the distance $D_B - D_S$. The cone of prices that support Pareto improving trades is shown as the two line segments emanating from A. Figure 10 (b) shows the situation if trades continue at the same price until the allocation reaches D_S. At that point, excess demand is the same as it was when the allocation was A, but individual incentives are much altered. Agent S has no more incentive to trade at the current price, and the cone of prices that support Pareto improving trades has narrowed considerably. Even when prices change over the course of a trading period, the change to excess demand is minor relative to the change in the cone of Pareto improving prices. So models based on excess demand neglect aspects of the exchange environment that will most likely prove to be important as incentive based models of price adjustment are developed. Models of this process, once developed, will almost certainly enhance further our understanding of the important problem of price dynamics in exchange economies.

References

[1] Anderson, Christopher M., Sander Granat, Charles R. Plott, and Ken-Ichi Shimomura (2004). "Global Instability in Experimental General Equilibrium: The Scarf Example," *Journal of Economic Theory*, **115**: 209 – 249.

[2] Arrow, Kenneth J. and Frank Hahn (1971). General Competitive Analysis. Amsterdam: North-Holland.

[3] Cox, James S. (1997). "On Testing the Utility Hypothesis," *Economic Journal*, **107**: 1054 – 1078.

[4] Davidson, Russell and James G. MacKinnon (1993). Estimation and Inference in Econometrics. Oxford: Oxford University Press.

[5] Easley, David and John O. Ledyard (1993). "Theories of Price Formation and Exchange in Double Oral Auctions," in Friedman, D., Rust, J. (Eds.), The Double Auction Market: Institutions, Theories, and Evidence, 63 – 97, Reading, MA: Addison-Wesley.

[6] Elliott, Graham, Thomas J. Rothenberg, and James H. Stock (1996). "Efficient Tests for an Autoregressive Unit Root," *Econometrica*, **64**, 813 – 836.

[7] Enders, Walter (1995). Applied Econometric Time Series. New York: Wiley.

[8] Friedman, Daniel (1991). "A Simple Testable Model of Price Formation in the Double Auction Market." *Journal of Economic Behavior and Organization*, **15**: 47 – 70.

[9] Gjerstad, Steven (2007). "The Competitive Market Paradox," *Journal of Economic Dynamics and Control*, **31**: 1753 – 1780.

[10] Gjerstad, Steven and John Dickhaut (1998). "Price Formation in Double Auctions," *Games and Economic Behavior*, **22**: 1 – 29.

[11] Gjerstad, Steven and Jason Shachat (2007). "Individual Rationality and Market Efficiency," Purdue University Economics Department Working Paper 1204.

[12] Hahn, Frank H. and Takashi Negishi (1962). "A Theorem on Non-tatonnement Stability," *Econometrica*, **30**: 463 – 469.

[13] Lian, Peng and Charles R. Plott (1998). "General equilibrium, markets, macroeconomics and money in a laboratory experimental environment," *Economic Theory*, **12**: 21 – 75.

[14] Scarf, Herbert (1960). "Some Examples of Global Instability of the Competitive Equilibrium," *International Economic Review*, **1**: 157 – 172.

[15] Smith, Vernon L. (1962). "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*, **LXX**: 111 – 137.

[16] ——— (1976). "Experimental Economics: Induced Value Theory," *American Economic Review*, **66**: 274 – 279.

[17] ——— (1982). "Microeconomic Systems as an Experimental Science," *American Economic Review*, **72**: 923 – 955.

[18] Williams, Arlington W., Vernon L. Smith, John O. Ledyard, and Steven Gjerstad (2000). "Concurrent Trading in Two Experimental Markets with Demand Interdependence," *Economic Theory*, **16**: 511 – 528. Reprinted in Studies in Economic Theory: Advances in Experimental Markets, Springer-Verlag, 2001.

[19] Wilson, Robert B. (1987). "On Equilibria of Bid-Ask Markets," in Feiwel, G.W. (Ed.), Arrow and the Ascent of Modern Economic Theory, 375 – 414. New York: New York University Press.

Appendix A

Prior to the start of an experiment session, a seller goes through an instruction set that describes each of the elements of the seller screen and their operations. These elements include: (1) the seller's endowment and current allocation of currency and commodity, (2) the seller's profit (or utility) function, (3) input that the seller provides during the session, (4) processing of input (by the CDA mechanism) from the seller, as well as from other sellers and from buyers, and (5) determination of the seller's profit, which is based on the seller's allocation of currency and commodity. Buyers' instructions are analogous to sellers' instructions. The most important principle adhered to in our implementation of instructions is that inputs of asks by sellers (or bids by buyers) are initiated by the seller (or by the buyer) in order to avoid creation of a reference point effect that would influence asks or bids once the session begins.

Each seller and each buyer views a total of 33 screens, and makes inputs on five of these screens. There are eight interactive questions that each seller and each buyer must answer correctly in order to proceed. Although there are 33 screens, most of the screen space is devoted to the actual screen display that the subject views and uses during the experiment session. During the instructions, a text box describes elements of the seller's screen display. The total length of the sellers' instructions in these text boxes is 3250 words, which is equivalent to approximately six pages of text. Subjects typically complete instructions in 15 to 35 minutes.

The sequence of steps through the screens is described below. The instruction summary below refers frequently to elements of the seller's screen, which is shown in figure 1 (p. 6). The instructions for a buyer are similar. Direct experimenter interaction with subjects was kept to a minimum whenever possible, including sign-in, seating, and payment.

Screen 1: The subject's earnings are based on subject's decisions and the decisions of other participants.

Screen 2: The subject is a seller throughout an experiment session that lasts for 12 trading periods (in the first three sessions) or 15 periods (in session Ariz1).

Screen 3: The subject should not communicate with or distract others. The subject's data are anonymous.

Screen 4: Payment is made at an exchange rate of £0.01 per unit of experiment currency accumulated. (This rate was $0.008 in session 4 at the University of Arizona.) Payment is made anonymously in cash at the conclusion of the experiment session.

Screen 5: There is a set of interactive instructions that follow this screen. The seller will know that he either has or has not completed instructions based on the status message at the top of his screen.

Screen 6: Each period of the session consists of a 60 second "Preview Phase", a "Trading Phase" of 180 seconds and a "Review Phase" of 30 seconds. A clock at the top of the screen ticks down to the end of each phase.

Screen 7: The seller begins each period with eighteen units of the commodity. The current balance of both currency and commodity are shown throughout each trading period in the Current Allocation box on the seller's screen.

Screen 8: The location and purpose of the "Profit Calculator" is described to the seller.

Screen 9: The seller is prompted to enter a Quantity and a Price into the Profit Calculator. The profit that would result if the seller sold the proposed quantity at the proposed price is shown in the Profit Calculator table, and the Current Allocation and Proposed Allocation are shown in the Profit Calculator Graph. (See figure 2 on p. 7 for a view of the Profit Calculator for a seller and for a buyer.)

Screen 10: The Profit Calculator table is described to the seller. Rows correspond to different prices and columns correspond to different numbers of units sold. The profit level for the price and quantity that the seller entered on Screen 9 is displayed in the center of the table.

Question 1: For the quantity 'q' and the price 'p' that the seller entered, he is asked to state what his profit would be if he were to sell '$q+1$' units at price 'p'.

Screen 11: The Profit Calculator Graph is described. The graph shows the seller's Current Allocation, the seller's Proposed Allocation (which is the allocation that would result if the seller exchanges q units of the commodity for $p\,q$ units of the numeraire commodity), and the Iso-profit Curve that passes through the Proposed Allocation.

Screen 12: The seller is reminded that he begins each period with eighteen units of the commodity, and is asked to use the Profit Calculator to determine what his total profit would be if he were to sell all eighteen units of the commodity at the price that he entered on Screen 9.

Question 2: The seller is asked to enter the profit that he would obtain if he were to sell all 18 units from his commodity endowment at the price that he entered on screen 9.

Screen 13: The Profit Calculator includes left and right arrows that can be used to decrease or increase the quantity proposal and up and down arrows that can be used to decrease or increase the price proposal. (See figures 1 and 2.) The seller is encouraged to test these features.

Screen 14: An ask is entered in the "Enter Ask" area of the seller screen. The subject enters an ask at this point.

Screen 15: The ask entered by the subject appears in the Unit Ask row of the Trade Summary display. Commodity Holding remains unchanged, but Commodity Available is reduced by one unit. The ask also appears in the "Your Asks" display area.

Screen 16: Asks by all sellers and bids by all buyers appear in the "Market Queue" display.

Screen 17: A new ask is generated randomly to simulate an action by another seller. The new ask is at a higher amount than the subject's ask, to illustrate the Ask Improvement rule. (The new ask is generated randomly from one to ten above the seller's own ask, though the subject is not informed of the process that generates the new ask.) The subject is also informed that if he makes a new ask at this point, the new ask will replace his current ask.

Question 3: The seller has a quiz question appear at this point. The seller, after seeing a description of the Ask Improvement rule, is prompted to state what is the highest ask that can be submitted at this point.

Question 4: After the subject answers Question 1 correctly, he is asked to state the number of asks that will be in the Market Queue if he submits a new ask.

Screen 18: The subject is asked to enter a new ask that improves on the ask submitted by the simulated seller.

Screen 19: The new ask replaces the seller's previous ask. The "Messages" display is shown, and updates to the Trade Summary display are described.

Screen 20: The subject is prompted to remove his ask by double clicking on the ask in the "Your Asks" display.

Screen 21: Changes to the seller screen that result from the removal of his ask are described, including updates to the Commodity Available, Your Asks, Market Queue, Trade Summary, and Messages displays.

Screen 22: The seller is prompted to enter a new ask to replace the ask that he just removed.

Screen 23: Updates to the seller's screen that result from the new ask are reviewed.

Screen 24: A trade occurs when a seller's ask is at or below the current best bid, or a buyer's bid meets or exceeds the best ask. The seller is informed that on the next screen, a bid from a simulated buyer will be generated that will meet the seller's ask, so that the seller will trade with the simulated buyer.

Screen 25: Updates to the seller's Unit Price, Unit Profit, and Total Profit entries in his Trade Summary table that result from the most recent trade are described. Changes to the Currency Holdings and Commodity Holdings areas are described. The Market Transaction Prices graph update is described.

Screen 26: The price determination rule is reviewed and the seller is informed that after two questions, a bid will be simulated that results in a trade with the current low ask.

Question 5: The seller is asked whether a trade will result if a buyer now submits a bid that is below the current low ask.

Question 6: The subject is asked what the trade price will be if a bid is submitted that meets or exceeds the current low ask.

Screen 27: A bid is simulated that generates another trade. (This trade is between a simulated seller and a simulated buyer, so the subject only sees public information regarding the trade, i.e., its price on the Market Transaction Prices graph.)

Screen 28: A new simulated bid appears in the Market Queue. The price determination rule is reviewed once more and the subject is asked two more questions (Questions 7 and 8).

Question 7: The seller is asked whether a trade will result if he submits an ask that exceeds the current high bid.

Question 8: The seller is asked what the trade price will be if he submits an ask that is below the current high bid.

Screen 29: The seller is prompted to enter an ask that is at or below the current high bid, in order to produce a new trade.

Screen 30: Changes to the Current Allocation, Trade Summary, and other screen displays that result from the most recent trade are reviewed.

Screen 31: The seller is informed of the 'Vote to End Period' option, and the unanimity rule that triggers an early end to the trading phase of the current period.

Screen 32: A Period Profit window appears during the review phase of each period.

Screen 33: Subjects are cautioned that the ask by another seller and the bids by other buyers in the instructions were simulated and that these may not be similar to the responses by buyers and by other sellers during the experiment. The subject is informed that he has now completed the instructions and trading will begin when all subjects have completed their instructions.

Appendix B

Starting from the adjustment process model

$$\Delta\hat{d}_{t+1} = -\gamma\,\hat{d}_t + u_{t+1}, \tag{B.1}$$

if errors are generated by the $AR(1)$ process $u_{t+1} = \rho\,u_t + \epsilon_{t+1}$, then substitution of u_{t+1} into equation (B.1) results in the equation

$$\Delta\hat{d}_{t+1} = -\gamma\,\hat{d}_t + \rho\,u_t + \epsilon_{t+1}. \tag{B.2}$$

Equation (B.1) for period t (rather than for period $t+1$) is $\Delta\hat{d}_t = -\gamma\,\hat{d}_{t-1} + u_t$. Solve this for u_t to get $u_t = \Delta\hat{d}_t + \gamma\,\hat{d}_{t-1}$. When this is substituted into equation (B.2), the result is

$$\Delta\hat{d}_{t+1} = -\gamma\,\hat{d}_t + \rho\,(\Delta\hat{d}_t + \gamma\,\hat{d}_{t-1}) + \epsilon_{t+1}. \tag{B.3}$$

This equation, with some minor rearrangement, is equivalent to an equation with lagged price changes on the right hand side and i.i.d. errors:

$$
\begin{aligned}
\Delta d_{t+1} &= -\gamma\,\hat{d}_t + \rho\,(\Delta\hat{d}_t + \gamma\,\hat{d}_{t-1}) + \epsilon_{t+1} \\
&= -\gamma\,\hat{d}_t + \rho\,\gamma\,\hat{d}_{t-1} + \rho\,\Delta\hat{d}_t + \epsilon_{t+1} \\
&= -\gamma\,\hat{d}_t + \rho\,\gamma\,\hat{d}_t - \rho\,\gamma\,\hat{d}_t + \rho\,\gamma\,\hat{d}_{t-1} + \rho\,\Delta\hat{d}_t + \epsilon_{t+1} \\
&= -\gamma\,(1-\rho)\,\hat{d}_t - \rho\,\gamma\,\Delta\hat{d}_t + \rho\,\Delta\hat{d}_t + \epsilon_{t+1} \\
&= -\gamma\,(1-\rho)\,\hat{d}_t - \rho\,(\gamma-1)\,\Delta\hat{d}_t + \epsilon_{t+1}\,.
\end{aligned}
$$