Days Without End: On the Stability of Experimental Single-Period Continuous Double Auction Markets

Dave Cliff, Chris Preist
Information Infrastructure Laboratory
HP Laboratories Bristol
HPL-2001-325
December 14th , 2001*

E-mail: cliffd@hplb.hpl.hp.com, cwp@hplb.hpl.hp.com

experimental economics, continuous double auction, equilibration, market dynamics, price stability

We present results from four experiments in which a simple continuous double auction market is operated on a network of computer terminals. Our experiments are based on those reported by Smith & Williams in a 1983 paper, but with one significant procedural difference: unlike in other experimental markets with which we are familiar, in our market there is no division of time into a sequence of discrete trading periods or "days" where, at the start of each period, all traders receive their full entitlement to buy or sell units of commodity, and the traders gradually exercise these until they are all simultaneously replenished at the start of the next trading period, at which point all unexercised entitlements from the previous trading period expire. Rather, in our experiments, traders' entitlements to buy and sell are updated periodically but the updates are sequenced out of phase in order to give a finer-grain discretization of time; and unexercised entitlements do not expire, but accrue until the end of the experiment. This change of method was introduced so that our experiments would have no direct correlate of discrete trading periods, and thus we could explore whether having discrete trading periods affects the dynamics of the market. Analysis of our results indicates that our single-period "continuous-time" markets showed equilibration dynamics remarkably similar to the "discrete-time" results presented by Smith & Williams. Thus, we conclude that our results offer no evidence to support the belief that the discretization of time in experimental markets affects the dynamics of those markets.

* Internal Accession Date Only
Approved for External Publication

1 Introduction

The continuous double auction (CDA) market mechanism has evolved over the years to play an increasingly important role in the world economy. International commodities markets, stock exchanges and futures market all use variants of the CDA as their mechanism for price determination. In recent years, many of these have moved to electronic trading, where traders input their bids and offers into a computer network, and a central computer determines which trades take place according to the market rules.

Because of the importance of the CDA market mechanism, much research has been carried out into the properties of it and its variants/refinements. This work has been both theoretical [2, 3] and experimental [6, 7, 9]. The theoretical work has tended to focus on a static formulation of the problem; it analyses the behaviour of a CDA marketplace in which each participant knows at the outset how many units of a good they are willing to trade, and the cost/value of each unit.

Smith [4] carried out pioneering experiments which consisted of a series of such static auctions; time was divided into 'days' or 'trading periods'. At the beginning of each trading period, all participants were informed of how many units they could buy or sell, and the value or cost of each unit. Smith showed that, even with a small number of participants, such a market would converge to equilibrium after only a small number of trading periods if the supply and demand remained constant. These results, and subsequent work [5] shows how robust the CDA institution is.

However, there is a difference between the experimental set-up used by Smith, and the workings of many of the CDA's used in practice. In the latter, trading cannot be divided into neat periods. In Smith's set-up, supply and demand is initialized at the beginning of a period. A trader cannot receive additional units to buy/sell, with accompanying values/costs, in the middle of a period. At the end of a period, any unsatisfied supply or demand disappears. It does not carry over into the next period. In most real-world CDA's, such as the stock markets, participants can receive more units to buy or sell at any time. Furthermore, there is no fixed "market-closing" deadline time by which everyone must trade or lose the opportunity to do so altogether. If a trader has failed to trade by the end of the day in a stock market, they can simply trade at the beginning of the next day if they so choose.

How important is this difference? Do impressive experimental results on the robustness and stability of the CDA with a small number of participants carry over to more realistic forms of CDA, ones not split into trading periods? Or does the end-of-trading-period play a role in bringing rapid convergence to equilibrium? In this paper, we give experimental results which strongly suggest that the division of time into trading periods does not have a significant effect on the stabilization of the marketplace.

2 Market Mechanism

To enable comparison with existing results we have endeavoured to make both our trading mechanism and experimental set-up as close as possible to the one described by Smith & Williams in a 1983 paper [8]. Smith & Williams experimented with various alternate rules in the CDA. One rule involved the use of a rank queue, in which bids (offers) are placed in a bid queue (offer queue), ordered from highest to lowest (lowest to highest). Each participant can have only one queued bid or offer; if they enter another, this replaces the original one. Traders may also withdraw their own bids or offers if they so choose.

The participant holding the best bid or offer cannot withdraw it, but can replace it with a better one. A trade takes place when the best bid or best offer is accepted. At this point, it is removed from the queue and the second bid/offer becomes the best. Smith & Williams showed that use of the rank queue resulted in more rapid convergence to equilibrium than a CDA without it. For that reason, we use this mechanism in the work presented in the paper.

In our experiments it is possible for a trader to have many units available for trade at a given time. For that reason, we have modified the rank queue to allow bids or offers for multiple units.

Buyers and sellers submit 'quotes' via individual computer terminals. Each quote lists the price and quantity a trader is willing to buy or sell. Quotes are added to the bid or offer queue using the same procedure as in Smith & Williams' experiment. However, one minor difference between our procedures and Smith & Williams' is that, in our experiments, traders accept an existing quote not by pressing a separate button, but by entering a new quote with a price that "crosses" the price of the existing quote. That is, to accept an existing bid quoted at price p_{bid}, a seller issues an offer with a price $p_{\text{off}} \leq p_{\text{bid}}$; similarly, if a buyer wishes to accept an existing offer quoted at price p_{off}, the buyer issues a bid priced at $p_{\text{bid}} \geq p_{\text{off}}$. In either event, this results in the highest bid-price being greater than or equal to the lowest offer-price, in which case a a trade takes place.

If the prices of the highest bid and offer cross but are different, trade takes place at the price entered first. If the quantities of the crossed quotes are different, then the quote with the smaller quantity is removed from the queue. The quote with the larger quantity is left on the queue but with its quantity reduced by the number of units traded. By using this mechanism, rather than an accept button, it is possible for traders to accept a quote, but trade a smaller quantity. In some marketplaces, this may not be appropriate: if selling part of a lot reduces the chance of selling the rest at a good price, a different mechanism should be used. However, in our experiments, this is not the case.

Each trader worked at a personal computer "terminal"; all terminals were connected to a network that allowed two-way communication with a central "server" computer. The server machine sent out data to be displayed on the traders' terminal screens, received data (e.g. bids and offers) from the traders, and performed all the book-keeping resulting from transactions and permit-updates.

All traders have the bid and offer queues displayed on their screen, and the price and quantity of the last 10 trades made in the market. This is the only public information.

Traders were informed that they would be issued with "trading permits" during the experiment that allowed them to buy or sell units of an imaginary commodity. Each permit entitled the holder to enter into a price-constrained transaction for one unit of the commodity: a single seller's permit would specify a particular transaction price-limit below which it could not be exercised; a single buyer's permit would specify a particular transaction price-limit above which it could not be exercised. The server computer rejected any quote issued by a trader who did not hold the necessary number of permits with appropriate price-limits that would be required if the quote was accepted in full by another trader. The limit prices on the permits correspond to the sellers' costs and the buyers' values in the traditional terminology of experimental economics. By phrasing the experiment in terms of permits, we obviated the need for explicitly endowing the buyers with money during the experiment: in effect, each buyer has infinite cash reserves but, because cash cannot be spent without a permit, it is the permits that induce the experimental market's demand schedule.

Rather than have all permits issued to a trader carry the same price limit, traders were told that they would be issued with permits carrying three different price-limits. We referred to these as the three different "types" of permit that any trader would be dealing with, and made it publically known that although every trader would deal with only three types of permit, there were more than three types of permit in the market so different traders may be allocated different combinations of permit-types. We also made it publically known that any trader could accumulate permits over the course of the experiment, and that unexercised permits would expire worthless at the end of the experiment. Each permit can be considered to be an order from an end-customer for one unit of the commodity at a given price, or a request from an end-customer to the trader to sell one unit fo the same commodity at a given price.

All traders had a "cash account" initialized to a zero balance. Traders were told that, for each permit exercised in a transaction at price p_{tra}, the seller's account would be credited with a "profit" of $p_{\mathrm{tra}} - p_{\mathrm{SL}}$ where p_{SL} is the price-limit on the permit that the seller exercises to enter into the transaction, and the buyer's account would be credited with a "profit" of $p_{\mathrm{BL}} - p_{\mathrm{tra}}$ where p_{BL} is the price-limit on the permit that the buyer exercises to enter into the transaction. When traders enter into a transaction, they do not need to specify which permit they use. The server automatically removes the permit which will make the maximum profit, and updates the profit information on the screen appropriately.

Each trader's screen display screen informs them of the price-limits of their three permit types, and the number of permits of this type they have in their posession. It also tells them the price of their last trade of each unit type, the profit they made on it, and the total profit they have made so far.

While each experiment was running, the server produced a log-file that recorded each bid, offer, withdrawal, permit-update, and transaction. For each event, the log-file recorded the time of the event (measured in seconds since the start of the experiment), an event-identification code-letter, and data such as the identity-number(s) of the trader(s) involved in the event, the cost or value of permit-updates, and prices of transactions. Informal observations indicated that the server's record of time could drift slightly from real-time, but never by more than a few (< 10) seconds over the course of a 44-minute experiment.

3 Experiment Design

The structural parameters used in the experiments are taken from Smith & Williams' 1983 paper[8]. The price-limits on each of the three types of permit issued to each trader are shown in Table 1, where the prices are defined as offsets from the competitive equilibrium price P_0. However, unlike Smith & Williams, we organised our experiments so that permits to buy or sell units are not all simultaneously allocated at the beginning of a trading period. Instead, permits enter the market at a pre-determined continuous rate, and are used up only when a trade is made.

	Type 1	Type 2	Type 3
Buyer 1	+0.95	−0.10	−0.25
Buyer 2	+0.70	0.00	−0.05
Buyer 3	+0.45	0.00	−0.15
Buyer 4	+0.25	+0.05	−0.20
Seller 1	−0.95	+0.10	+0.25
Seller 2	−0.70	0.00	+0.05
Seller 3	−0.45	0.00	+0.15
Seller 4	−0.25	−0.05	+0.20

Table 1: Smith & Williams' schedule of traders' permit price-limits (i.e., buyers' values and sellers' costs), expressed as deviation from the competitive equilibrium price P_0.

The trading takes place continuously for a period of 44 minutes. Once every 120 seconds, each player receives one Type 1 permit, one Type 2 permit, and one Type 3 permit. However, these 120-second update cycles are "staggered" (i.e., out of phase) for all players: no player receives an update of more more than one permit; and all players have to wait at least 10 seconds between updates, with some players having considerably longer inter-update periods. Full details of the staggered update schedule is shown in Table 2.

Comparison of Tables 1 and 2 reveals that, within any 120-second time-window, all eight traders will each receive three new permits (one of each type). Hence a 120 second time-window in our experiments is analagous to a single trading period in Smith & Williams' experiments, except that Smith & Williams re-supplied all 24 entitlements to buy and sell at the start of each trading period, whereas in our experiments the permits enter the market in pairs at 10-second intervals.

Time	+0	+10	+20	+30	+40	+50	+60	+70	+80	+90	+100	+110
Buyer permit	+95	+70	+45	+25	+5	0	0	-5	-10	-15	-20	-25
Seller permit	-95	-70	-45	-25	-5	0	0	+5	+10	+15	+20	+25

Table 2: Times (seconds into the cycle) and values (expressed as P_0 offsets, in cents) of permit updates

This means that at time t=0 seconds, Buyer 1 and Seller 1 each receive one permit to trade a unit of Type 1. At t=10, Buyer 2 and Seller 2 each receive a permit to trade a unit of Type 1, and so on. Then, at t=120 the cycle repeats, with Buyer 1 and Seller 1 each receiving another Type 1 permit, regardless of whether they have exercised the Type 1 permits issued to them 120 seconds earlier.

Note that whenever a buyer receives a permit to buy a unit of a given value $P_0 + \Delta$, a seller will simultaneously receive a permit to sell a unit of cost $P_0 - \Delta$. Because of this it is theoretically always possible for any participant receiving a permit for an inframarginal unit to exercise it immediately. Furthermore, provided no extramarginal units have been involved in trade, the equilibrium price at any given time (assuming the market cleared instantly at that time) would be P_0, or a price range including P_0.

Our subjects were undergraduate and graduate students at Stanford University, recruited via email and poster advertising. No subjects had participated in previous CDA experiments, and were not permitted to participate more than once. They were paid $15 on arrival, and allocated a terminal and a role (from Buyer 1 to Seller 4). Subjects were issued with a set of instructions to read and used a tutorial program to familiarize themselves with how to enter and withdraw quotes. At the end of the experiment, subjects were paid in private, according to their performance. However, the different limit prices on each trader's different types of permit implied that, under equilibrium conditions, different traders could be expected to earn final profits differing by a factor of up to about three. Thus, to equalize the potential real financial gain of all subjects, the subjects were each privately supplied with an individual "exchange rate" for converting from their profit balances to the (US Dollar) currency we paid them in, such that in equilibrium conditions all traders would be expected to earn the same amount. The expected payment to each subject was approximately $40.

4 Results

Four experiments were conducted. For brevity, will refer to them using the labels E1, E2, E3, and E4. In each experiment there were eight traders: four buyers and four sellers.

The experiments reported here were the first four continuous-time experiments ever conducted on our experimental economics facility: a network of trading terminals that had been developed primarily to study markets where trading happened in discrete periods, and where traders' submission of quotes or other market events were comparatively infrequent events. Unfortunately, our real-time experiments reported here placed an exceptionally heavy load on the server and on the network that allowed the server to communicate with the traders' terminals. In E1 and E4, the network loading became excessive and response-times (e.g., the delay between submitting a bid and that bid appearing on the bid-queues of all traders' screens) grew unacceptably long, prompting us to prematurely terminate those two experiments. E1 was terminated after 30 minutes, while E4 was terminated after 22 minutes. In both cases, the traders ceased to submit quotes, but the

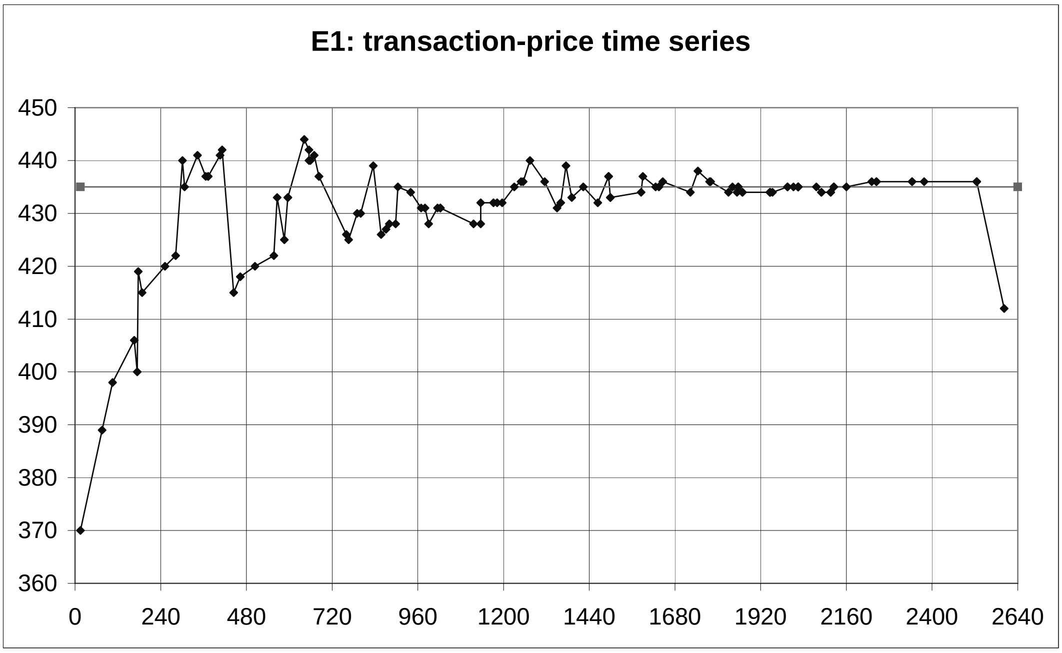

Figure 1: Transaction-price time series for E1.

backlog of quotes already submitted continued to be processed by the central server, so transactions continued to occur after the traders had been told to cease trading. This fact should be considered when interpreting the data shown below for E1 and E4.

Transaction-price time-series from our four experiments are shown in Figures 1 to 4. In each figure, the horizontal axis shows time as recorded by the server, while the vertical axis shows transaction prices. For each experiment, the theoretical competitive equilibrium price (denoted by P_0) is illustrated by a horizontal line, with a small square marker at the point where the horizontal line intersects the vertical (price) axis.

In traditional period-based experiments, the deviation of transaction prices from P_0 is very often illustrated by a metric based on the root mean squared (RMS) difference between equilibrium and transaction prices during each trading period – Smith (1963) denoted this metric as σ_0, while Smith & Williams (1983) used the symbol α to denote an identical metric.

Following Smith (1963), we use $\sigma_0(p)$ to denote the RMS price deviation for a trading period p. If there are n_p transactions in period p, then $\sigma_0(p)$ is defined as:

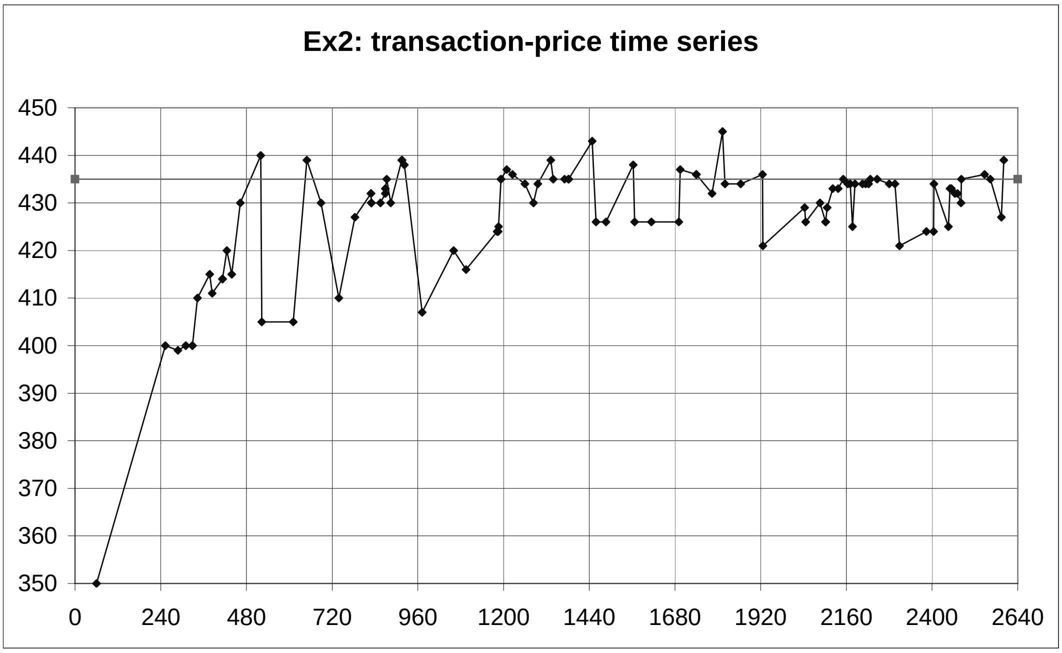

Figure 2: Transaction-price time series for E2.

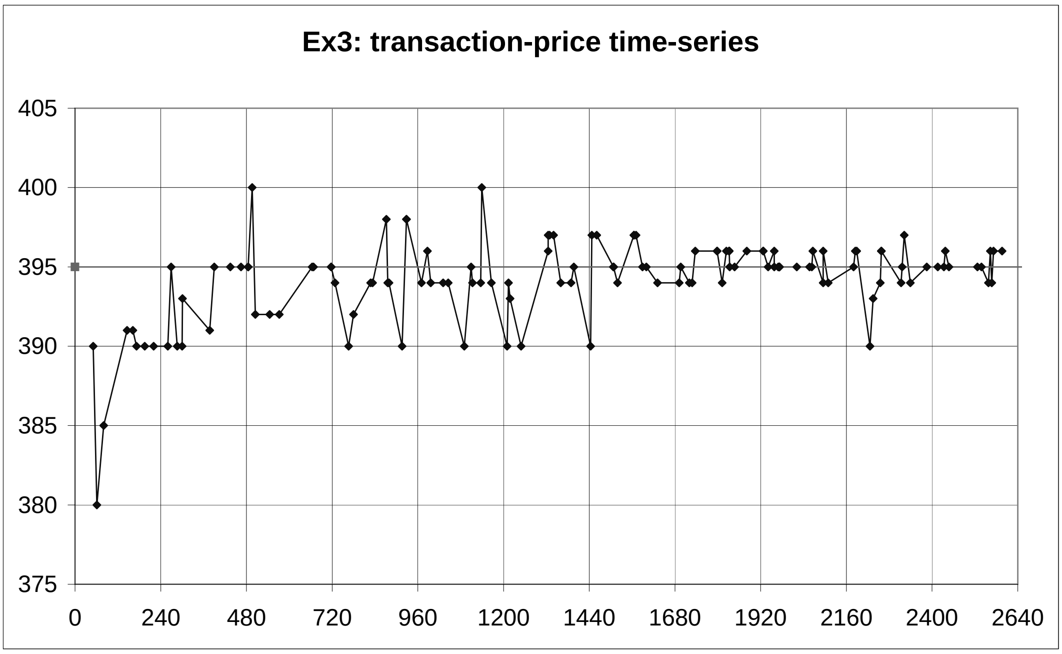

Figure 3: Transaction-price time series for E3.

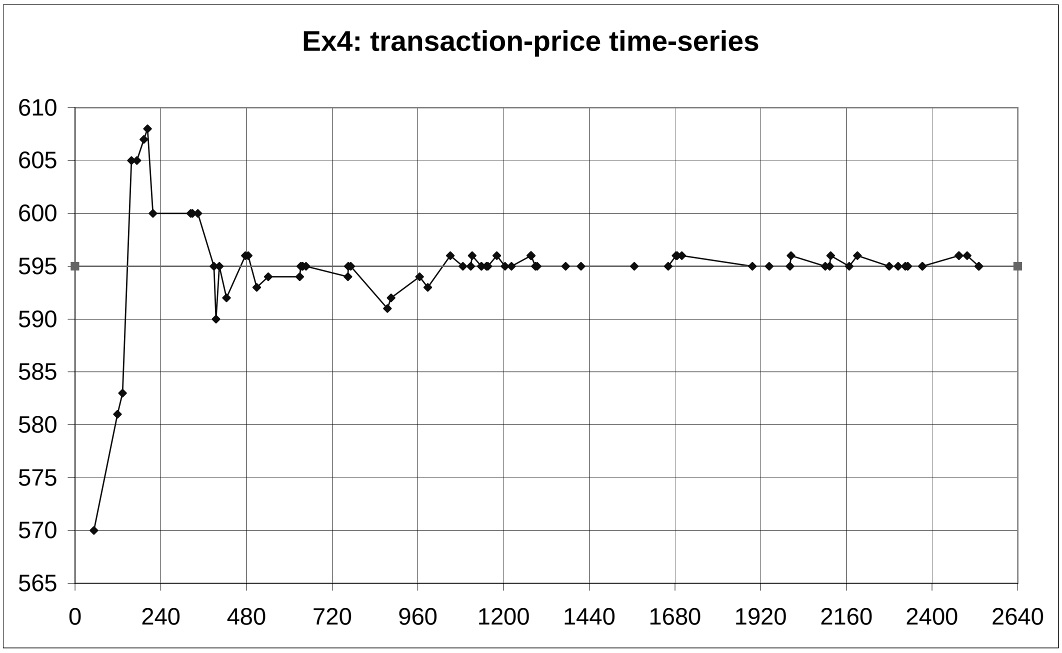

Figure 4: Transaction-price time series for E4.

$$\sigma_0(p) = \sqrt{\frac{1}{n_p}\sum_{i=1}^{n_p}(p_i - P_0)^2} \tag{1}$$

Where $p_{i:i\in 1,\ldots,n_p}$ are the prices of the transactions in period p, and P_0 is the theoretical competitive equilibrium price.

While such RMS-based metrics have a natural appeal, the period-based definition requires extension for continuous-time data. Recall that in our experiments the distribution of new permits is scheduled in such a way that, in any 120-second time-window, the entire supply and demand arrays of Smith & Williams (1983) are (re-)allocated to the traders. For this reason, we use a 120-second window in calculating a RMS deviation metric at time t. That is, we consider all transactions that have occurred over the preceding 120 seconds. Using $\mathcal{T}_P[t_1, t_2]$ to denote the set of prices of all transactions occuring between times t_1 and t_2, we denote our RMS deviation metric as $\sigma_{120}(t)$ where:

$$\sigma_{120}(t) = \sqrt{\frac{1}{|\mathcal{T}_P[t_s, t]|}\sum_{p_j\in\mathcal{T}_P[t_s,t]}(p_j - P_0)^2} \tag{2}$$

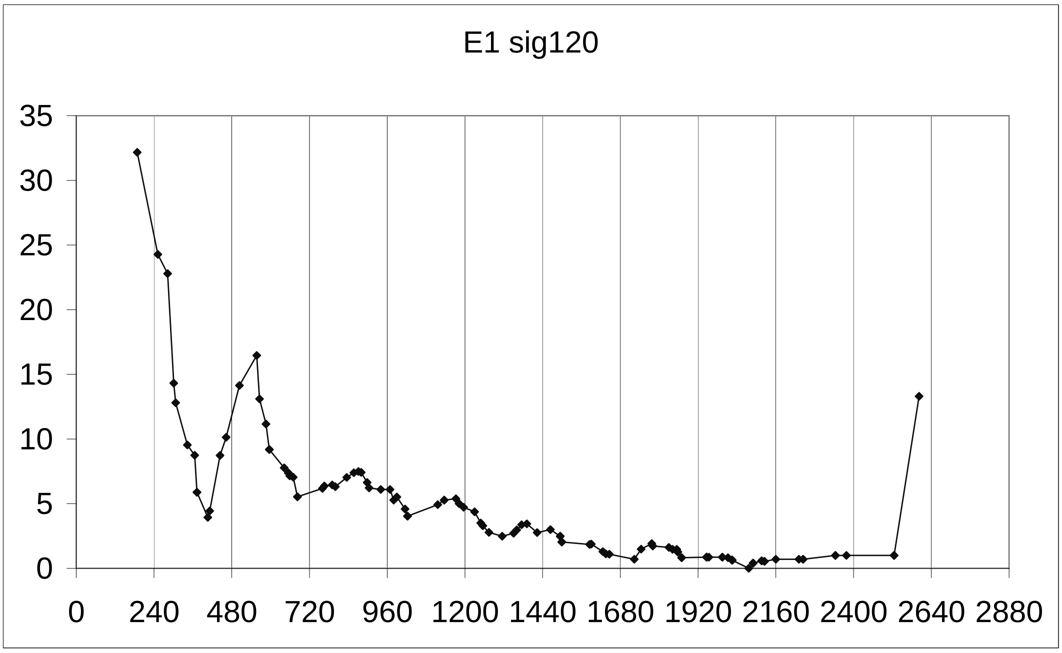

Figure 5: Time-series of $\sigma_{120}(t)$ for E1.

and

$$t_s = \begin{cases} t - 120 & \text{if } t \geq 120 \\ 0 & \text{otherwise} \end{cases}$$

The corresponding $\sigma_{120}(t)$ time-series for each experiment is illustrated in Figures 5 to 8. As the figures clearly show, in all four experiments the initial RMS deviations from transaction prices are high, but they rapidly fall in the first few minutes of the experiments.

To allow closer comparison, Figure 9 shows all four σ_{120} time-series. It is clear from Figure 9 that in three of the experiments (E1, E3, and E4), the σ_{120} values fall to below 0.10 within 600 seconds, and fall to below 0.05 by t=1200s. But in one experiment (E2) the values of σ_{120} over the same range of times are consistently higher: although the σ_{120} time series does fall below 0.10 around t=840s, it subsequently rises again and does not return to below 0.10 again until around t=1200s; for the remaining 1400s until the end of the experiment, σ_{120} values vary between 0.10 and 0.05. Investigation of the E2 log-file revealed that this difference in the σ_{120} data was due

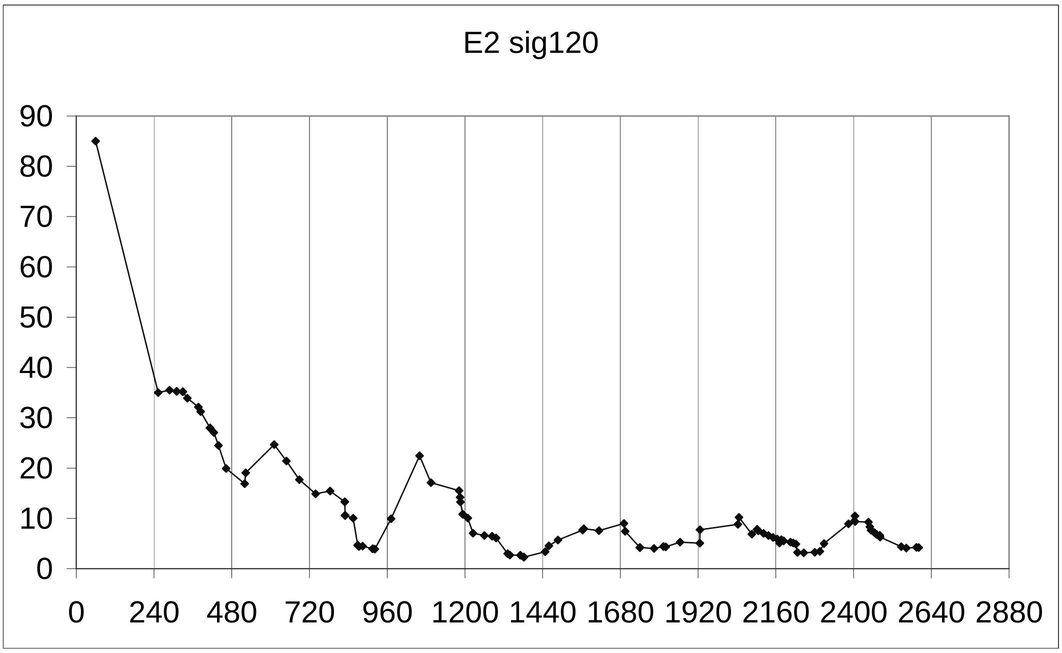

Figure 6: Time-series of $\sigma_{120}(t)$ for E2.

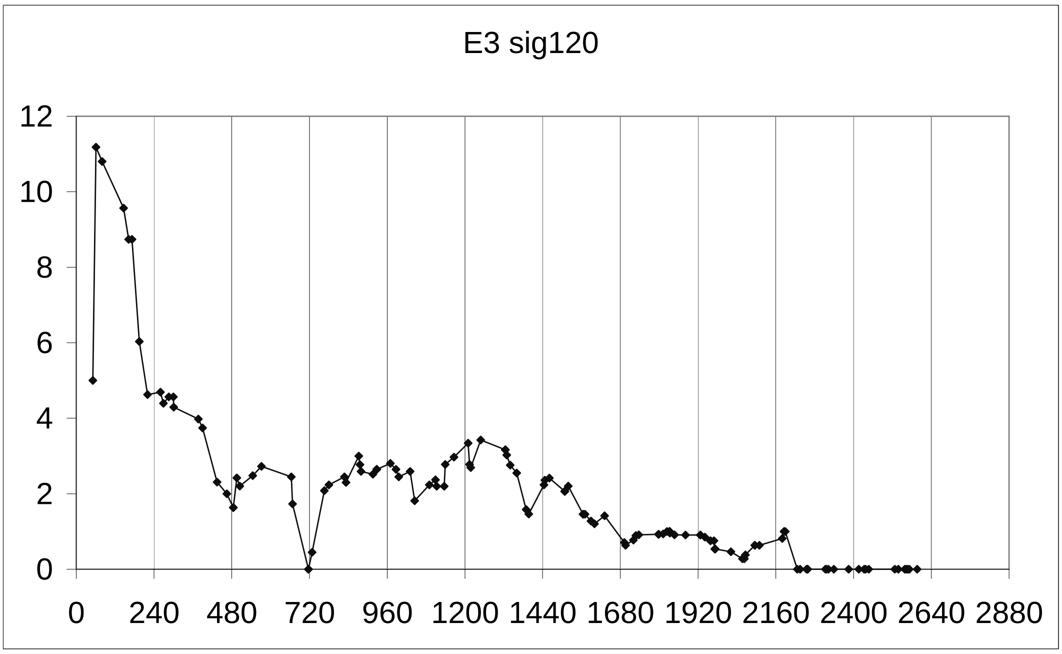

Figure 7: Time-series of $\sigma_{120}(t)$ for E3.

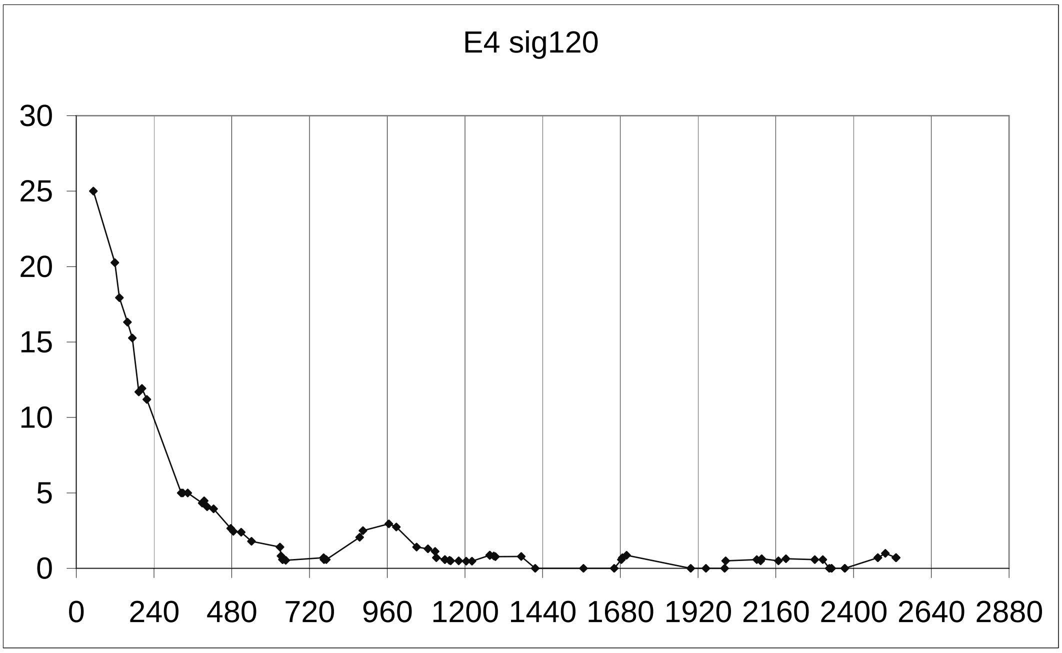

Figure 8: Time-series of $\sigma_{120}(t)$ for E4.

to the presence of a "soft seller": one of the sellers was consistently entering into transactions at prices significantly lower than could have been achieved if that seller had been more "aggressive", holding out for a higher price. The activity of this one soft seller was sufficient to alter the σ_{120} data for the entire experiment.

5 Discussion

In their 1983 paper, Smith & Williams present a regression analysis of their data from period-based experimental CDAs with experienced and inexperienced subjects. They assume an exponential relationship between trading period index p and their deviation measure $\sigma_0(p)$ (Equation 1), defined as:

$$\sigma_0(p) = e^{a+bp+cX_s} \tag{3}$$

where a is a constant, b is a coefficient weighting the period-index p, and c is a coefficient that weights

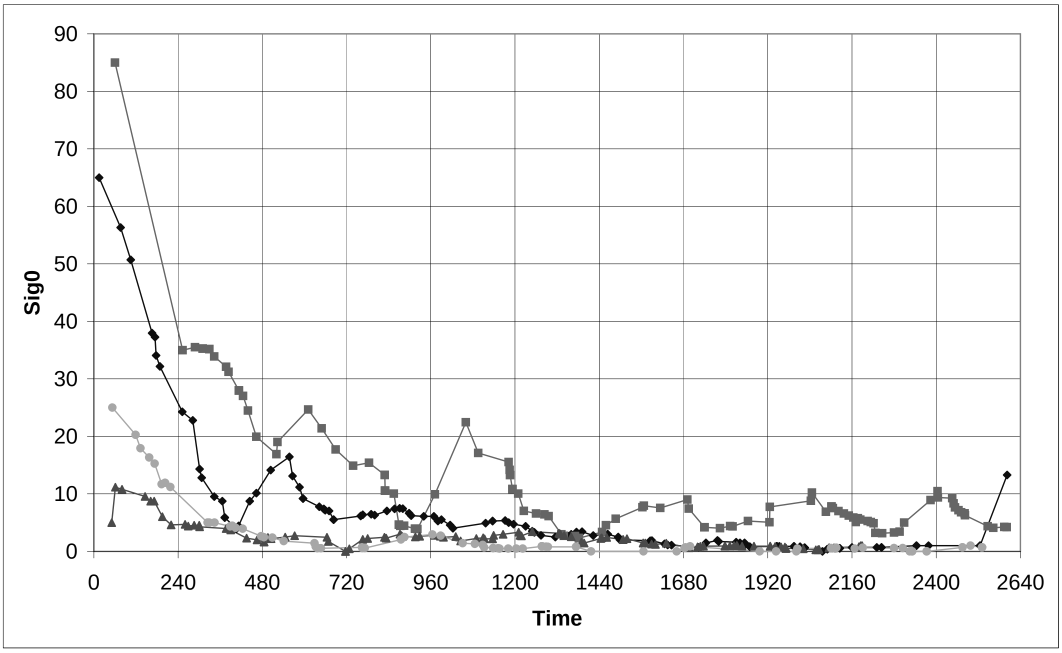

Figure 9: Time-series of $\sigma_{120}(t)$ for E1 to E4.

a subject-experience variable X_s: X_s is binary, equal to one if the subjects in the experiment have prior experience of experimental CDA markets, and equal to zero if the subjects are inexperienced.

Smith & Williams conducted three experiments comparable to ours, i.e. using inexperienced subjects in a CDA market with ranked order queues for bids and offers. They show one transaction-price time series from such an experiment (1983, p.318), and then summarise their data by fitting curves of the form given in Equation 3 for $X_s = 0$ and $X_s = 1$ to the data of all three experiments. Smith & Williams state that the following parameter-values were obtained by regression: $a = -1.609$, $b = -0.205$, $c = -0.872$.

To compare the data from our non-iterated experiments with Smith & Williams' data, some minor manipulations of our data are necessary. The aim of these manipulations is to compensate for the lack of discrete trading periods in our experiments by partitioning our time-series data into "pseudo-periods", thereby enabling the use of Equation 3's discrete period-based functional form.

As was discussed earlier, the entire supply and demand arrays used by Smith & Williams flow into our market once every 120 seconds. Thus, we use time windows of 120 seconds as the basis for calculating our discretized data: we divide time into 120-second "pseudo-periods", indexed by $\tilde{p} : \tilde{p} \in \{1, 2, 3, \ldots\}$. Using $\mathcal{T}_T[t_1, t_2]$ to denote the set of times at which transactions occur in the period between time t_1 and time t_2, we define a simple arithmetic mean statistic $\widehat{\sigma}_{120}(\tilde{p})$:

$$\widehat{\sigma}_{120}(\tilde{p}) = \frac{1}{n_{\tilde{p}}} \sum_{i=1}^{n_{\tilde{p}}} \sigma_{120}(t_i) \tag{4}$$

where $n_{\tilde{p}} = |\mathcal{T}_T[120(\tilde{p} - 1), 120\tilde{p}]|$ and $t_{i:i \in \{1,2,\ldots,n_{\tilde{p}}\}} \in \mathcal{T}_T[120(\tilde{p} - 1), 120\tilde{p}]$.

Figure 10 summarises the $\widehat{\sigma}_{120}$ data generated from the transaction-price time-series for experiments E1 to E4, illustrating the mean and standard deviation of the four $\widehat{\sigma}_{120}$ series. As can be seen, our 44-minute experiments yield 22 pseudo-periods.

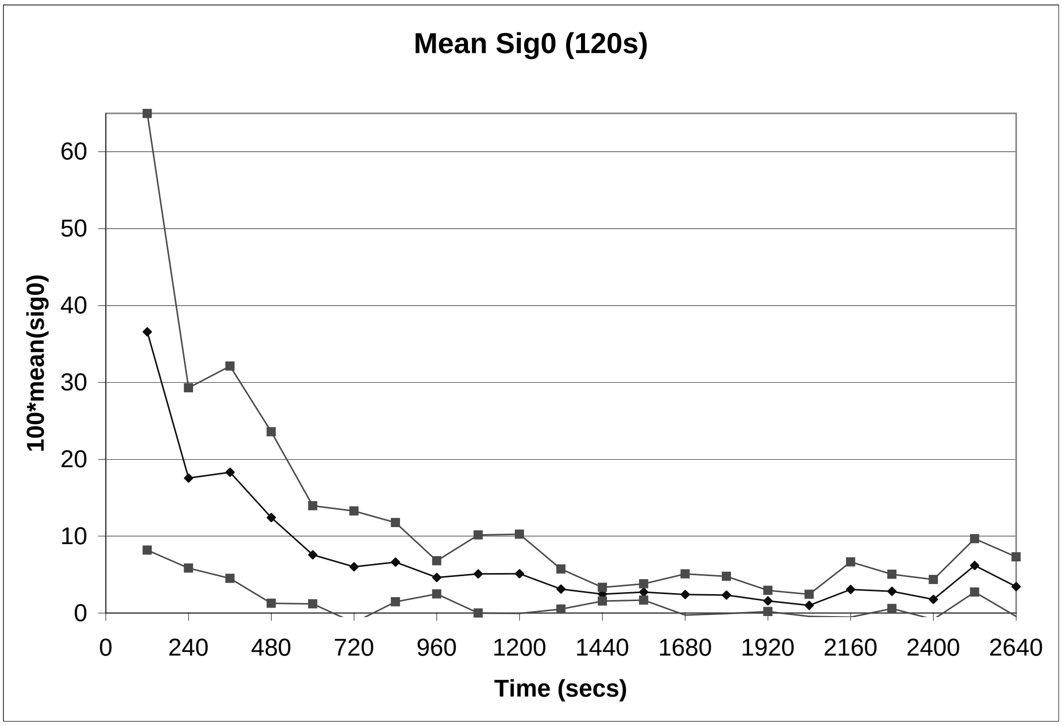

Figure 10: Discretized $\widehat{\sigma}_{120}$ data: mean, plus and minus one standard deviation.

We then performed regression analysis on the mean discretized data shown in Figure 10. Because all subjects in our experiments were inexperienced, $X_s = 0$ and hence the value of c for our data is irrelevant. Thus we can derive values corresponding to Smith & Williams' a and b, for our $\widehat{\sigma}_{120}$

data using linear regression to find values of $\tilde{a}$ and $\tilde{b}$ for:

$$\ln(\widehat{\sigma}_{120}(\tilde{p})) = \tilde{a} + \tilde{b}\tilde{p}$$

Our data series of 22 pseudo-periods is almost three times longer than the 8-period data presented by Smith & Williams (1983). To allow direct comparison, we only considered values of $\tilde{\sigma}_{120}(\tilde{p})$ over the range $\tilde{p} = 1$ to $\tilde{p} = 8$. Regression yielded values of $\tilde{a} = -1.51324$ and $\tilde{b} = -0.21036$. In percentage terms, our value of $\tilde{a}$ differs from Smith & Williams' a by approximately 3% and our value of $\tilde{b}$ differs from Smith & Williams' b by approximately 6%.

The difference between the curve generated by Smith & Williams' data and the curve our discretized data generates is shown in Figure 11, which for comparison also shows the mean $\widehat{\sigma}_{120}(\tilde{p})$ data from Figure 10.

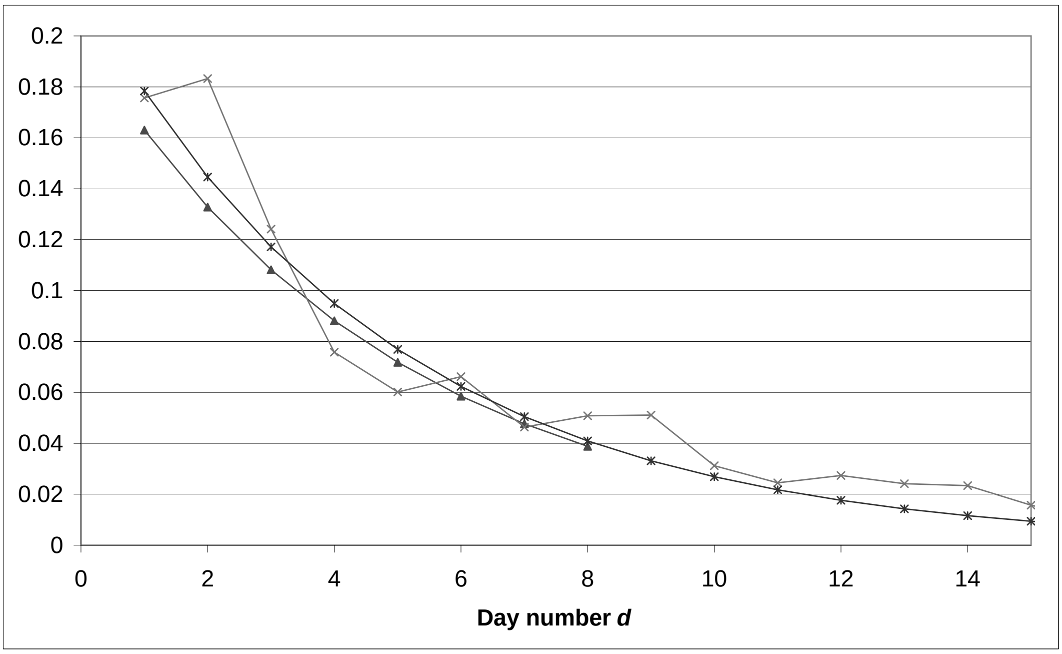

Figure 11: Curves obtained by regression: Smith & Williams' (1983) curve (triangles); curve for our discretized data (asterisks); mean $\widehat{\sigma}_{120}$ data replicated from Figure 10 (crosses).

There are two notable features in Figure 11. First, our curve is so close to Smith & Williams' that it is practically indistinguishable – especially when the standard deviations shown in Figure 10 are taken into consideration. Second, although our curve is given by values of $\tilde{a}$ and $\tilde{b}$ derived by regression over the range $\tilde{p} = 1$ to $\tilde{p} = 8$, when the curve is continued for values $\tilde{p} > 8$ it remains close to the mean $\widehat{\sigma}_{120}$ data. On the basis of these observations, we conclude that the curve generated from our values $\tilde{a}$ and $\tilde{b}$ is a good summary representation of our entire data, and offers no reason to suspect that the data from our experiments differs from those reported by Smith & Williams (1983).

6 Conclusion

At the time of conducting these experiments (Summer 1998) we had been unable to locate any published descriptions of non-period-based CDA markets in the experimental economics literature. Thus, to the best of our knowledge, the four experiments reported in this paper are the first continuous-time CDA markets studied under laboratory conditions. Subsequently, Brewer *et al.* [1] published a technical report in December 1999 which explored continuously-refreshed supply and demand environments.

As was made clear in the discussion, there is no principled reason for concluding that our data differ significantly from that presented by Smith & Williams (1983). Our experiments were designed with the intention that the only significant difference in methods was the replacement of the traditional sequential discrete trading periods with a single 'period' lasting the duration of the entire experiment. Because our data are in such good agreement with Smith & Williams' data, we conclude that the transaction-price time-series of their period-based CDA markets do not differ significantly from those of our "continuous time" CDA markets. This result invites a number of possible lines of future research, the most pressing of which is investigating the equilibration behavior of markets where the stock equilibrium differs from the flow equilibrium.

Acknowledgements

Many thanks to Steve Gjerstad for his cooperation, wisdom, guidance, hospitality and good humour while we conducted these experiments in the Hewlett-Packard Labs Palo Alto (HPL-PA) experimental economics facility. Thanks also to Jason Shachat for valuable discussions and for his help in running the experiments, and to Shailendra Jain for hosting us in his department at HPL-PA. Thanks also to Lin Jones for her help in the preparation of this paper.

References

[1] P. J. Brewer, M. Huang, B. Nelson, & C. R. Plott (1999). *On the Behavioral Foundations of the Law of Supply and Demand: Human Convergence and Robot Randomness.* Social Science Paper no.1079, Pasadena, California: California Institute of Technology.

[2] A. Rustichini, M. Satterthwaite, & S. Williams (1994). Convergence to Efficiency in a Simple Market with Incomplete Information. *Econometrica* 62(5):1041–1063, 1994.

[3] M. Satterthwaite & S. Williams (1989). Bilateral Trade with the Sealed Bid k-Double Auction: Existence and Efficiency. *Journal of Economic Theory* 48:107–133, 1989.

[4] V. L. Smith (1962). An Expeirmental Study of Competitive Market Behavior. *Journal of Political Economy*, 70:111–137, 1962.

[5] V. L. Smith (1992) *Papers in Experimental Economics.* Cambridge: Cambridge University Press, 1992.

[6] V. L. Smith (1982). Microeconomic Systems as an Experimental Science. *American Economic Review* 72(5):923–955, 1982.

[7] V. L. Smith, G. L. Suchanek, & A. W. Williams (1988). Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets. *Econometrica* 56(5):1119–1151, 1988.

[8] V. L. Smith & A. W. Williams (1983). An experimental comparison of alternative rules for competitive market exchange. In M. Engelbrecht-Wiggans et al. (editors) *Auctions, Bidding, and Contracting: Uses and Theory.* New York: New York University Press, 1983. pp.307–334.

[9] V. L. Smith & A. W. Williams (1990). The Boundaries of Competitive Price Theory: Convergence, Expectations, and Transaction Costs. in L. Green and J. H. Kagel (editors) *Advances in Behavioral Economics*, Ablex Publishing, 1990.