Sunspots in the Laboratory

John Duffy*
Department of Economics
University of Pittsburgh
4S01 Posvar Hall, 230 S. Bouquet Street
Pittsburgh, PA 15260, USA
+1-412-648-1733
jduffy@pitt.edu

and

Eric O'N. Fisher
Department of Economics
The Ohio State University
410 Arps Hall, 1945 North High Street
Columbus, OH 43210, USA
+1-614-292-2009
fisher.244@osu.edu

First Draft: 28 October 2001
This Draft: March 28, 2002

Abstract

This paper shows that extrinsic uncertainty influences markets in a controlled environment. It provides the first direct evidence of "sunspot" equilibria. Sunspots in the laboratory can be quite sensitive to the flow of information. Extrinsic uncertainty matters when information flows slowly as in a closed book call market, but it need not matter when there is a fast flow of information, as in a double auction, where infra-marginal bids and offers are observable.

Keywords: Sunspots, Extrinsic Uncertainty, Experimental Economics
JEL Classification: D5, C9

*The authors gratefully acknowledge funding from the National Science Foundation under grant numbers SES-0111123 and SES-0111315. They also Jack Ochs, Tim Cason, and seminar participants at the fall 2001 Economic Science Association Meetings and at the University of Iowa for helpful comments on earlier drafts.

1. Introduction

The effects of extrinsic uncertainty have fascinated social scientists since well before Mackay (1841). Is there some kind of randomness, having nothing to do with fundamentals, that nevertheless serves as a way of coordinating the expectations and consequent plans of market participants? This question has received sound theoretical foundations in the analysis of Cass and Shell (1983) and Azariadis (1981), and it has spawned a vast literature on "sunspot" equilibria in macroeconomics, finance, and other fields.[1]

Despite the extensive theoretical attention that has been paid to sunspot equilibria, there is little *direct* evidence showing that sunspot variable realizations are responsible for any of the economic volatility observed in actual markets.[2] The difficulty lies in identifying sunspot variables and isolating their effects from those of fundamental variables on market activity. A possible resolution to this problem is to examine market behavior in a controlled environment where realizations of sunspot variables can be isolated from shocks to preferences or endowments. In this paper, we describe such an experimental design and our findings. In particular, we report direct evidence--the first ever--of the existence of sunspot equilibria.

We examine sunspot equilibria in "call markets" where a market-clearing price is determined in an environment akin to a Walrasian system. We show that sunspot

[1] Farmer (1999) gives a very good summary of the importance of these ideas in macroeconomics. Jevons (1884) used the term "sunspot" because he mistakenly believed that solar activity drove the business cycle. In the modern parlance, a sunspot is any random variable that is unrelated to fundamental factors, like endowments, preferences, or technology.

[2] There is quite a lot of *indirect* evidence using calibrated general equilibrium business cycle models, which exploit the possibility that the set of equilibria in such models may be indeterminate. This indeterminacy allows self-fulfilling beliefs or sunspot variable realizations to become an additional source of volatility in these models. See Benhabib and Farmer (1999) for a survey of this literature.

equilibria are possible in these markets; indeed, they *always* occur in this treatment of our experiment. We also find that sunspot equilibria are less likely to obtain in environments where markets clear in a more decentralized way, namely in a computerized double auction. Our findings thus suggest that sunspot equilibria are less likely to arise in decentralized environments. This is an important finding as it casts doubt on the usefulness of sunspot equilibria as an explanation for volatility in markets where prices are not determined according to a centralized market clearing mechanism. Thus, sunspots may be more likely to arise on the NYSE, where a market maker sets price according to his order book, than on the NASDAQ, an over-the-counter market where it is common to see several outstanding bids and offers at any one time.

2. Related Literature

We are not the first to use the laboratory in an effort to obtain direct evidence of sunspot equilibria. Marimon, Spear, and Sunder (1993) report results from an experiment designed to implement an economy with overlapping generations where sunspots may play a role. Under the assumption of perfect foresight, their environment has two steady-state equilibria and one where prices follow a two-period cycle. This multiplicity allows for the possibility that prices depend upon a sunspot variable.

Marimon, Spear, and Sunder tried to use realizations of a sunspot variable to coordinate expectations on the cyclic equilibrium. Their variable consisted of a blinking cube on the computer screen that alternated in color between red and yellow. In the absence of any correlation between sunspot realizations and actual price movements, they found that subjects ignored the sunspot variable and simply coordinated on one of the two steady states. Consequently, they sought to induce a correlation between price movements and sunspot realizations in each session's "training" periods by alternating the number of subjects assigned to play the role of "young" agents. This design amounted to an endowment shock that was perfectly correlated with realizations of the sunspot variable, and it did induce a cycle in prices. Once the training period was over, the shock to economic fundamentals was eliminated. Marimon, Spear, and Sunder found that prices fluctuated according to the sunspot realization in the training phase. But once that initial period ended, they found in most sessions that price volatility quickly dampened and subjects coordinated near one of the two non-cyclic steady-state equilibria. In those sessions where prices remained volatile after the training period, the actual price path deviated substantially from the predicted two-period cyclic equilibrium. Thus, while a significant effort was made to get subjects to condition their expectations on a sunspot variable, Marimon, Spear, and Sunder did not observe a sunspot equilibrium in any of their five sessions.

Our design differs considerably from theirs. In particular, we consider a simpler, partial equilibrium framework that allows us to abstract from a number of conceptual difficulties. In our simple and static environment, there are two equilibria, and market-clearing quantities do not differ across equilibria. Our sunspot variable consists of an announcement by the experimenter about the likely state of the market. The announcement has no bearing on the true state of the market, which depends instead on the decisions of the subjects themselves, as explained below. The announcement simply serves as a potential coordination device that subjects may use or ignore.

We believe that our sunspot--the announcement of the likely state of the market--provides the necessary additional context that was missing from Marimon, Sunder, and

Spear's blinking cube. Still, our announcement is a genuine sunspot variable, since it has nothing to do with economic fundamentals. Furthermore, unlike those authors, we are able to obtain coordination on sunspot equilibria without resorting to any stochastic fundamentals. Indeed, they argue that conditioning on such randomness is necessary for generating sunspot-induced volatility.[3]

There is some related experimental work involving games with multiple equilibria where experimenters have examined how subjects respond to *recommendations* by the experimenter as to how to play the game. For examples, see Brandts and Holt (1992), Brandts and MacLeod (1995), and Van Huyck, Gillette, and Battalio (1992). While the aim of this literature is different from the our goals, one finding that emerges is that pairs of subjects will follow a recommendation as long as it neither involves play of a dominated strategy nor results in asymmetric payoffs. By contrast, in the markets we consider, we find that subjects are willing to coordinate on our announcements even though some subjects strictly prefer one state of the world to the other. This finding likely obtains because an individual subject has less influence in a market than in a two-person games.[4]

Finally, we note that there is a relationship between the sunspot equilibria we examine and the notion of a *correlated equilibrium*.[5] The realizations of our sunspot variable give subjects a means of implementing a self-enforcing correlated strategy. The announcement of the likely state of the world provides a common signal about which

---

[3] Marimon, Spear and Sunder (1993, p. 77) state, "Before these cyclic movements can be supported solely by extrinsic signals (or sunspots) subjects *must be* exposed to intrinsic events that are correlated with the extrinsic variables." (We added the emphasis.)

[4] In a recent experiment that is somewhat similar to ours, Ball, Eckel, and Zame (2001) show that arbitrary assignments of subjects into "high status" and "low status" groups can affect prices in a game with multiple equilibria; in particular, they find that prices are higher when high status sellers meet low status buyers.

state will actually occur, and each subject takes an action conditional upon that common signal. Our subjects submit bids or offers that depend upon their subjective beliefs about the states that may occur. In our call market treatment, unilateral deviations from this correlated equilibrium are unprofitable, and the sunspot equilibrium can also be regarded as a correlated Nash equilibrium of a simple game. To our knowledge, there are no laboratory investigations of correlated equilibrium, so our findings should also be of interest to those working with this concept.

3. Hypotheses

There are two fundamental hypotheses that we explore. The first is:

HYPOTHESIS 1: *It is possible to construct a laboratory environment in which sunspots matter. Further, these sunspot equilibria can be easily replicated.*

As we show below, coordination on a sunspot equilibrium *always* obtains in the call market, and occasionally obtains in the double auction market. Our findings suggest that coordination on a sunspot equilibrium is easily replicated, a result that will allow others to build upon our design. While the logical foundations of equilibrium theory based upon endogenous expectations of intrinsic uncertainty are quite well founded, the econometric evidence based on data from field markets is mixed at best. Indeed, Flood and Garber's seminal work (1980) showed how difficult it is to find sound evidence for price bubbles using econometric tests based upon a well-specified model. Hence, there is compelling need for evidence from the laboratory.

Our second hypothesis is subtler, and it is perhaps of greatest interest to both economic theorists and policy makers.

---

[5] See Aumann (1974) for the definition of correlated equilibrium, and Peck and Shell (1991) on the relationship between correlated and sunspot equilibria.

HYPOTHESIS 2: *Sunspot equilibria can be sensitive to the flow of information*.

The usual Walrasian framework that serves as the foundation for any theory of extrinsic uncertainty is based upon a static notion of the flow of information. It actually obviates an important element of many field markets, where there is nearly continuous trading between events that signal the advent of significant new information. The simplest way to allow for differential flow of information in asset markets in the laboratory is to highlight the difference between a double auction, in which several transactions prices can occur in a period, and a call market, in which by design only one price clears the market in each period. Of course, in a double auction, all the infra-marginal bids and offers become a part of the information set of every trader as the period unfolds, while only the marginal bid and marginal offer become public knowledge in a call market at the end of the period, once the price fixing has occurred.

4. Experimental Design

We conducted a 2×2 experimental design in which the treatment variables were the *market mechanism* and the *forecast announcement*. The two different markets are a double auction and a call market, and both were computerized.[6] The forecast announcement serves as our sunspot variable. Either the experimenter's random number generator predetermines the announcement, or a subject publicly flips a coins and makes an announcement that anyone can verify. The four cells of our experimental design are presented in Table 1.

Table 1

		Announcements	
		Pre-determined by Experimenter	Public Coin Flips
Market Mechanism	Double Auction	Cell 1 3 Sessions	Cell 2 3 Sessions
	Call Market	Cell 3 3 Sessions	Cell 4 3 Sessions

All four treatments are otherwise identical. In particular, there are always two equilibria, one with a low price and one with a high price, and the equilibrium quantity is identical. The equilibria are not Pareto comparable by design; some agents prefer one state or the other. If one particular equilibrium were Pareto dominant such as in a sunspot model of unemployment, then subjects might coordinate on it as a focal point for their expectations. On the other hand, if each equilibrium gave rise to the same infra-marginal rents, then sunspots would matter only in a trivial sense, since every subject's payoff would be independent of the state of nature.

Let $\tilde{p}$ be the median transaction price at the end of a trading period. In a call market, every transaction occurs at this price, but we are using this formalism to allow for a double auction as well. The median, end of period transaction price was chosen to mitigate the effect that any one transaction might have on the determination of $\tilde{p}$. In order to focus on the effects of extrinsic uncertainty, we state that the "high state" occurs if $\tilde{p} \in H$ and the "low state" occurs if $\tilde{p} \in L$. The sets H and L form a partition of R_+,

[6] It is customary in an experimental economics paper to include the instructions that subjects received. But this paper is already too long, so we have omitted them. We urge the interested reader to retrieve the instructions for any treatment at http://economics.sbs.ohio-state.edu/efisher/duffyfisher.

considered as the set of all non-negative prices. Thus $H \cap L = \varnothing$ and $H \cup L = R_+$. In practice, we set $L = [0,b)$ and $H = [b,\infty]$, where b is a cutoff price.

Each agent can purchase or sell up to two units. The marginal valuations for the $i-th$ buyer are $v_i(s) = (v_{i1}(s), v_{i2}(s))$, where the fact that these depend upon the state $s \in H \cup L$ is made quite explicit. Likewise, the marginal costs of the $j-th$ seller are $c_j(s) = (c_{j1}(s), c_{j2}(s))$. As usual, we impose that $v_{i1}(s) \geq v_{i2}(s)$ and $c_{j1}(s) \leq c_{j2}(s)$; thus these valuations give rise to state-dependent individual demand correspondences or firm supply correspondences.

The $i-th$ buyer's individual demand correspondence is:

$$d_i(p,s) = \begin{cases} 0 & if \;\; p > v_{1i}(s) \\ \{0,1\} & if \;\; p = v_{1i}(s) \\ 1 & if \;\; p \in (v_{1i}(s), v_{2i}(s)) \\ \{1,2\} & if \;\; p = v_{2i}(s) \\ 2 & if \;\; p < v_{2i}(s). \end{cases}$$

Likewise, the individual supply correspondences are given by:

$$s_j(p,s) = \begin{cases} 0 & if \;\; p < c_{1j}(s) \\ \{0,1\} & if \;\; p = c_{1j}(s) \\ 1 & if \;\; p \in (c_{1j}(s), c_{2j}(s)) \\ \{1,2\} & if \;\; p = c_{2j}(s) \\ 2 & if \;\; p > c_{2j}(s). \end{cases}$$

Market demand is then $D(p,s) = \sum_i d_i(p,s)$, and market supply is $S(p,s) = \sum_j s_j(p,s)$, and the excess demand correspondence is $Z(p,s) = D(p,s) - S(p,s)$. An equilibrium is a function $p(s)$ and corresponding quantities demanded $D(p,s)$ and supplied $S(p,s)$ such that $0 \in Z(p,s)$ for all $s \in H \cup L$.

Because the median price itself determines the state of nature, it is easy to construct treatments where there are actually two equilibria for this market. In our treatments, there are five agents on each side of market. Let $p_L < p_H$, and recall that $\tilde{p}$ is the median transaction price. The two equilibria have the property that

$$D(p_H, \tilde{p} \in \{H\}) = S(p_H, \tilde{p} \in \{H\}) = D(p_L, \tilde{p} \in \{L\}) = S(p_L, \tilde{p} \in \{L\}) = 6.$$

Thus the equilibrium *quantities* are independent of the state of nature. Again, we emphasize that these two equilibria are not Pareto comparable by design. Two buyers and two sellers do better when the state of nature is low, two buyers and sellers do better when the state of nature is high, and the remaining buyer and seller earn the same rents in both states.

Figure 1 shows the actual demand and supply curves that we used in the all experimental sessions. The steps are drawn to indicate precisely which valuations and costs accrue to which subjects. From the figure it is clear that $p_L \in [90,110]$ and $p_H \in [190, 210]$. In all the treatments, we set $L = [0,150)$ and $H = [150, \infty)$. Hence, if the median transaction price at the end of a period is less than 150, the state of nature is low; otherwise it is high. The most important fact about these parameters is that every subject had to make a decision based on uncertainty. Each subject had to make bids or offers not knowing which state would occur. The low or high costs or values might obtain, depending upon the realization of the end-of-period median price, as discussed below.

Let $\tilde{A}$ be a random variable whose support is $\{H, L\}$. A realization of this variable $a \in \{H, L\}$ corresponds to an announcement by the experimenter about the likely state of the market or the result of a coin toss where heads corresponds to state *H* and tails to state L[7]. A *sunspot* equilibrium is defined by the property $0 \in Z(p(a), a)$ for all $a \in \widetilde{A}$. *Sunspots matter* if $a \neq a'$ implies that $p(a) \neq p(a')$ because the agents' payoffs differ (by design) across realizations of the random variable $\widetilde{A}$.

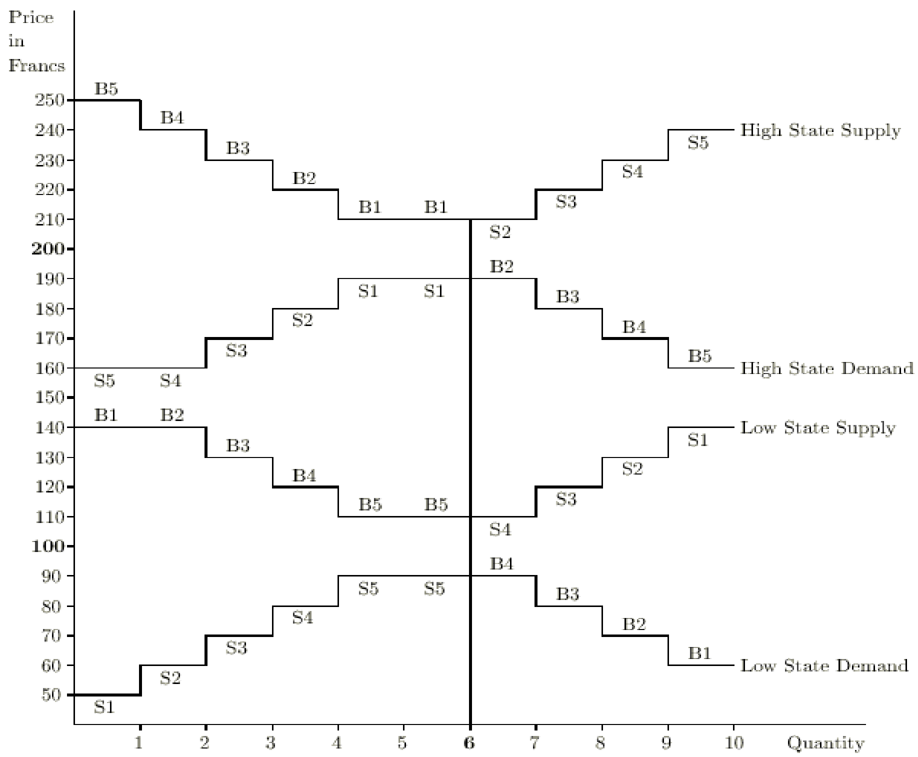

Figure 1: Induced Demand and Supply in High and Low States
Buyers: B1–B5, Sellers: S1–S5

During the first three periods of each session, we trained the subjects by eliminating the low state of nature. Thus the demanders had high valuations, the suppliers had high costs, and the equilibrium was supported by the price near $p_H = 200$. During the next three periods, we eliminated the high state of nature, and only the parameters for

[7] In the treatments, the announcements have the natural interpretation according to common language.

a low state were germane. Thus during the first six training periods in any session, the subjects learned how to use the computerized software while they were replicating two different static environments, one where $p_H \in [190, 210]$ was a supporting price and then one where $p_L \in [90,110]$ was likewise. It is well known that even in static and replicated environments in the laboratory, it takes several periods for the equilibrium to converge. That is in part why we used these first six training periods. We also wanted to make the equilibrium prices and quantities focal points for the subsequent periods in which extrinsic uncertainty was allowed full and free rein. Still, throughout the instructions, we made it very clear that the state of nature would be determined eventually by subjects' decisions alone.

Indeed, starting with period 7 and continuing until period 16, the state of nature was endogenous. We use data only from these ten periods to test our hypotheses. At the beginning of each of these periods, either the experimenter made a public announcement (based upon a predetermined random draw of random variable) or one subject (a different one every round) publicly flipped a coin to determine the announcement. It was understood that heads meant one should announce "high," and tails meant that one should announce "low." The exact phrasing of the announcement, in either case: "The forecast is high," or "The forecast is low."

It is important to give the exact the text of the relevant instructions. (These were read out loud and given in writing to each subject.) In the treatment where the experimenter made the announcement, they were:

Thus we would expect that $p_H = p(H) > p_L = p(L)$, although there is another equilibrium in which the price is perfectly negatively correlated with the announcement.

"Beginning with period 7, the experimenter will make an announcement at the beginning of each period. The announcement will be either that "the forecast is high" or that "the forecast is low." It is important that you understand that these announcements are only *forecasts*; they may be wrong, and they do not determine in any way your actual costs or values in that period. Indeed, the experimenter does not have any more information than you do. Remember that your actual costs and values depend only upon the official median price for that period."

In keeping with the spirit of the literature on sunspots, we used a random number generator to determine the announcements. Since the sequence of announcements is an obvious control variable, the same sequence was used in every session in Cells 1 and 3; utilizing the same sequence in several sessions allows us to determine whether a particular sunspot equilibrium can be replicated. Table 2 gives the specific sequence that we used in all "experimenter announcement" sessions.

Table 2

Round	7	8	9	10	11	12	13	14	15	16
Announcement	Low	High	Low	Low	High	High	Low	High	High	Low

In the treatment where the realization of a coin flip was used to determine the announcement, the instructions read:

"Beginning with period 7, an announcement will be made at the beginning of each period. The announcement will be either that "the forecast is high" or the "the forecast is low. This forecast will be determined by flipping a coin. Anyone who wants to can come up and look at the coin and how it landed. If the coin lands heads up, the person who flipped it will announce that the "forecast is high." If it lands tails up, that person will announce that the "forecast is low." The experimenter will ask each of you to take a turn flipping the coin. When it is your turn, flip the coin in the air and let it land on the floor. Anyone can come up at any time, and make sure that the person making the announcement is telling the truth. I will now let everyone see that this is a fair coin, and I will keep the coin in plain view at every moment during the experiment. Come up and look at the coin now."

In this treatment, the random sequence of announcements will necessarily differ across sessions, so that replication of a particular sequence of announcements is not possible in this treatment. Still, we can examine whether coordination on sequences of announcements obtains across all sessions in Cells 2 and 4, where coin flips occur.

We chose this treatment for two reasons. First, a public randomization device might matter. Second, our findings for the treatment where announcements were made by the experimenter might be subject to a Clever-Hans effect.[8] Specifically, we were concerned that subjects might place undue reliance on the experimenter's announcement because they were afraid that something nasty was in store if they tried to deviate from it. Alternatively, they may have blindly followed the experimenter's announcement because they wanted to please the experimenter or had trust in a professor. The coin-flip treatment allows us assess whether such Clever-Hans effects were present. In particular, it is of independent interest to see whether behavior in the two market environments differs when the stochastic process used to determine announcements is more transparent and more obviously beyond the control of the experimenter.

The instructions in both announcement treatments made it clear to subjects that announcements were not binding in any way. Indeed, subjects were reminded that their actual costs and values depended *only* upon the official (median) price for that period in the double auction or on the official market price for that period in the call market environment. Thus each trader was faced with a decision fraught with uncertainty about which state would actually occur,[9] and some of them learned by hard knocks that the

[8] This effect is discussed widely in experimental psychology. It captures the notion that subjects may respond unconsciously to cues that the experimenter is giving unwittingly. Indeed, in the nineteenth century ago, Clever Hans was a famous German horse who could do arithmetic by tapping out answers with his hoof. Rigorous investigation revealed eventually that he was responding to subtle (often unconscious) cues that his (human) audience gave him.

[9] In the double auction, the standard improvement rule was used, and the current median transaction price was always written in plain sight on the blackboard as transactions occurred during the period. In the call market, each subject submitted either two bids or two offers. The bids were ranked to make a demand

official market price could be quite different from what was announced of from what they had hoped would occur.

In the computerized double auction, subjects were allowed to submit bids or offers as long as they had units left to buy or sell. A trading period lasted for four minutes. Subjects observed the best bid and ask prices on their screens and could sell or buy at these prices. Bids and offers followed standard improvement rules; to get onto everyone's screen, a buyer has to increase the standing bid and a seller has to undercut the standing offer. When a transaction occurred, subjects saw the price at which the unit was exchanged, and all bids and asks on the computer screen were cleared. Also, the experimenter reported the current median traded price based on all transactions that had occurred in that period up to that point in time. The *official* median price that determines the actual values and costs that subjects faced in periods 7 thorough 6 was thus finally determined until the end of the four-minute trading period.

In the computerized call market, subjects typed in a positive integer for their first bid or offer and then typed in a second positive integer for their second bid or offer. The computer program then sorted all bids from highest to lowest and all offers from lowest to highest, thus creating demand and supply schedules $D(p)$ and $S(p)$. The market-clearing price $\tilde{p}$ was determined was follows. If there was an interval $[\underline{p}, \overline{p}]$ such that for all $p \in [\underline{p}, \overline{p}]$ $D(p) - S(p) = 0$, then price $\tilde{p}$ was the integer value closest to $\tilde{p} = (\underline{p} + \overline{p})/2$.[10] Sellers whose offers were less than or equal to $\tilde{p}$ sold all such units.

---

curve, and the offers were ranked to make a supply curve. These two curves determined the equilibrium, with the (almost standard) rule that splits the surplus equally between the marginal buyer and seller.

[10] If $p = \underline{p} = \overline{p}$ and $D(p) - S(p) \neq 0$, then a lottery was conducted among those on the long side of the market to determine who got to trade.

Likewise, buyers who bid at least $\widetilde{p}$ could purchase a unit for each such bid. Each subject was informed of the market-clearing price as well as the number of units bought or sold. The call market mechanism was carefully explained to subjects, using several illustrative examples.

In both the double auction and the call market, buyers' payoffs were the induced values net of transacted price, and sellers' payoffs were transacted price net of unit costs. Subjects could (and occasionally did) lose money if they bought too dear or sold cheap.

5. Experimental Results

We conducted twelve sessions, consisting of three sessions in each cell in Table 1. Each session involved 10 inexperienced subjects recruited from the undergraduate populations of the Ohio State University or the University of Pittsburgh. Two out of the three sessions in Cell 1 and two of the three sessions in Cell 2 were done in Columbus and the rest of the sessions in Cells 1-2 were done at Pittsburgh. We obtained the very similar findings using both subject pools, so we tentatively concluded that the subject pool was not important. The six call market sessions in Cells 3 and 4 were all done at the University of Pittsburgh.

The experiments produced two important results. First, equilibria where sunspots matter *can* be constructed and replicated in the laboratory; thus we find strong support for Hypothesis 1. Second, whether sunspots matter appears to be quite sensitive to the flow of information. Indeed, sunspot equilibria arise always only in the call market environment and only occasionally in the double auction. Since a call market has a much more restricted flow of information than a double auction, we conclude that there is support for Hypothesis 2.

A sunspot equilibrium obtains only if *every time* the announcement is high the median price $\tilde{p} \in [150, \infty)$ and *every time* the announcement is low the resulting median or market price lies in the range $\tilde{p} \in [0,150)$. Still, in judging whether a sunspot equilibrium obtains, we must allow for some noise in the experimental data. The design of the experiment predicts $D(p(a),a) = S(p(a),a) = 6$, but we often observed transaction volumes that were different from this. So we will still say that a *sunspot equilibrium obtains empirically* even if $D(p(a),a) = S(p(a),a) \neq 6$.

5.1 Double Auction, Experimenter Announcements

Figures 2, 3, and 4 show the data from the treatment in Cell 1, a double auction where the experimenter made the pre-determined announcements in Table 2. The figures show every transaction and the predicted prices $p_L = 100$ or $p_H = 200$. Market volume and median prices are not reported, but it they are easy to infer form the figures..

--Insert Fig. 2 here.--

--Insert Fig. 3 here.--

--Insert Fig. 4 here.--

These figures show that a sunspot equilibrium never occurred in these three sessions. While market volume is near the prediction of 6 units in every trading period, in the two sessions shown in Figures 2 and 4 the equilibrium differs from the announcement in at least two trading periods, so according to our (very stringent) criterion, a sunspot equilibrium does not obtain empirically. In Figure 3, the subjects coordinated near the low equilibrium in each of the final ten rounds; in this session, the sunspot announcement is completely ignored by subjects. .

We have an explanation for why subjects do not coordinate on sunspot equilibria in this treatment. In analyzing the first bids and offers in each period, we came to believe that demanders who benefited most in a high period tried to induce that state of nature by making a high opening bid; likewise suppliers who benefited most in a low period tried to get in the very first low offer. The standard improvement rule for a double auction then makes it impossible for another demander to bid lower even if he were fairly sure that a low state were likely to occur. The same fact is true for another seller who might be fairly sure that a high state will occur. Thus the flow of information in a double auction does seem to allow infra-marginal bids and offers to serve as signals independent of the sunspot realization. The initial transactions are very important, a fact that has been found in field data as well. Section 6 gives a formal model of information flow as a cascade.

5.2 Call Market, Experimenter Announcements

Figures 5, 6, and 7 show the data from the treatments with a call market where the experimenter again made the sequence of announcements described in Table 2. In this case, the figures plot the (single median) price as determined $\tilde{p}$ and transaction volume too. Again, these data are reported for just the last ten periods, those in which we are testing for sunspots. These figures contrast sharply with Figures 2 through 4; the evidence that subjects coordinated on the sunspot announcements is quite clear. An analysis of the infra-marginal bids and offers does seem to indicate that some subjects are submitting bids and offers in a strategic attempt to influence the state of nature, but this is a much more formidable task in a call market than in a double auction. Indeed, it takes about five or six (independent) bids or offers--each betting in essence against the sunspot announcement--to make the market clearing price move above of below the threshold that defines the state of nature. The only thing a subject knows is her own bids or offers and the market price, revealed at the end of the period. This paucity of information seems to make it too difficult to influence the state of nature in a strategic way, and the risk of making the wrong bid or offer is just too great to try to buck the trend that the sunspot announcement signals.

--Insert Fig. 5 here.--

--Insert Fig. 6 here.--

--Insert Fig. 7 here.--

5.3 Double Auction, Coin-Flip Determined Announcements

Figures 8, 9, and 10 show the data from the treatments with a computerized double auction, where announcements were determined by having each of the 10 subjects flipping a coin in turn. In this treatment and the next, the sequence of announcements varies across sessions; this variation is reflected in the differences in the predicted price sequence across sessions, which correspond precisely to the sequence of publicly and randomly determined announcements. It is quite interesting that we observe two sunspot equilibria (Figures 8 and 9) in this treatment. Although a sunspot equilibrium *is* possible in a double auction, there is no guarantee that it will occur. In the one session (Figure 10) where there was no sunspot equilibrium, the subjects coordinated on the low price equilibrium, a similar result to that found in second session of the double auction treatment with experimenter announcements (shown in Figure 3).

--Insert Fig. 8 here.--

--Insert Fig. 9 here.--

--Insert Fig. 10 here.--

What accounts for the difference in behavior between the double auction treatment with predetermined announcements and the treatment with the public randomization device? One might speculate that the greater transparency of the randomization device in the coin-flip treatment played a role. On the other hand, closer inspection reveals that the results for the double auction sessions are not all that dissimilar across the two announcement treatments. In two of the three double auction-experimenter announcement sessions, the subjects are *close* to coordinating on a sunspot equilibrium, though according to our criterion, coordination does not obtain. Furthermore, in both treatments there is a single session where the sunspot announcements are essentially ignored and coordination on the low price equilibrium obtains. We conclude that while coordination on sunspot equilibria in double auction environments is possible, such coordination may be quite fragile. In section 6 we provide a model that can account for this fragility.

5.4 Call Market, Coin-Flip Determined Announcements

Figures 11, 12, and 13 show the data from the treatments with a computerized call market, where announcements were determined by having each of the 10 subjects take a turn flipping a coin. The coin-flip determined announcements are evident from the predicted market price. For all three sessions, coordination on a sunspot equilibrium obtains. We conclude that coordination on sunspot equilibria in the call market environment is robust to the manner in which random announcements are made.

5.5 Discussion

The results presented in Figures 2 through 13 provide clear support for our hypotheses. Sunspot equilibria can be shown to exist in a controlled environment, and we are the first to have produced them. Indeed, we observe a sunspot equilibrium in 8 out of 12 sessions. In the call market, sunspot equilibria *always arise*. The latter finding is readily replicated, either in the strict sense, where the same sequence of sunspot realizations is used in different sessions with different subjects, or in the weaker sense where the environment remains the same, but the history of sunspot realizations differs owing to random coin flips. By contrast, sunspot realizations may not occur in the double auction environment, a finding that leads us to conclude that the occurrence of sunspot equilibria depends on the *market mechanism*. Table 3 supports conclusion. It summarizes our findings from all 12 sessions. Using the null hypothesis of a random assignment of successes (sunspot equilibria observed) across the two treatments, Fisher's exact test[11] has a p-value of 0.03. So the null hypothesis is easily rejected for a test of size 5%. We conclude that the market mechanism matters for producing sunspot equilibria. We now turn our attention to explaining *why* the market mechanism matters.

Table 3

		Sunspot	
		Was Observed	Not Observed
Market Mechanism	Double Auction	2	4
	Call Market	6	0

6. A Model of Information Flow in a Double Auction

This section develops a simple model of an information cascade that explains how initial transactions in the double auction environment can determine all the other transactions that follow in that period. The model is useful in understanding subjects'

[11] See, e.g. Siegel and Castellan (1988) for the precise details of this nonparametric test.

incentives to ignore or to follow realizations of the sunspot variable. Indeed, one clever subject who makes an early transaction can induce the state of the world that he prefers, regardless of what announcement has been made. This possibility arises only in the double auction environment, where trading and the flow of information are much more decentralized than in a call market. We believe this difference accounts for the disparate experimental outcomes reported in section 5.

The main risk a subject faces in rounds 7 through 16 is that his forecast of the state of nature might turn out to be wrong. A buyer who is fairly certain that the state will be high will make an early bid in a neighborhood of $p_H = 200$, thinking that his high valuations will obtain. Likewise, a seller who is somewhat sure that the low state will obtain is willing to make an early offer in a neighborhood of $p_L = 100$. There is an asymmetry inherent in the standard improvement rule in a double auction: once a buyer has put in a bid near the high price, no lower bid has standing until after a transaction has occurred. There is an analogous asymmetry on the supply side: a low offer trumps all higher ones.

There is also an asymmetry in the risks that buyers and sellers face. *Every* buyer would like to make an early transaction in a neighborhood of $p_L = 100$ because she will make normal profits if the state does turn out to be low and she will make extraordinary profits is the state actually turns out to be high. Likewise, *every* seller would like to make an early transaction in a neighborhood of $p_H = 200$ because a mistaken forecast about the state of nature can only redound to her benefit. Table 4 summarizes the relevant information.[12]

[12] The buyers are B1 through B5, and sellers are S1 through S5.

Table 4: Subjects' Infra-Marginal Rents
on First Unit, Second Unit (if traded)

	High State Transact at $p_H = 200$	Low State Transact at $p_L = 100$	Low State Transact at $p_H = 200$	High State Transact at $p_L = 100$
B1	10, 10	40	-60	110, 110
B2	20	40	-60	120
B3	30	30	-70	130
B4	40	20	-80	140
B5	50	10, 10	-90, -90	150
S1	10, 10	50	150	-90, -90
S2	20	40	140	-80
S3	30	30	130	-70
S4	40	20	120	-60
S5	40	10, 10	110, 110	-60

Several points are worth mentioning. First, the treatments were designed so that the high and low states were not Pareto comparable; in the terminology of Cass and Shell (1983), *sunspots matter* in these treatments. Thus there are some subjects—both on the demand side and the supply side—who strictly prefer a high state to a low one. Second, every, buyer would like to transact at a low price, and every seller would like to transact at a high price. These strategies are risk dominant in the sense that a buyer who actually transacts at $p_L = 100$ can never lose money and a seller who makes a sale at $p_H = 200$ will never lose money either.

6.1 An Information Cascade

A complete game-theoretic specification of a double auction in continuous time and with uncertain valuations is well beyond the scope of this paper. But there is a simple decision-theoretic model that explains the data that we observed. Assume, for simplicity, that the action set for each subject is restricted to $A_i = \{100, 200\}$. This is not an unrealistic assumption since the first six periods have been used to train the subjects in the markets without uncertainty, and the markets converge amazingly quickly to equilibrium in these early rounds. By the time the seventh period begins, each player knows he can make at least one profitable transaction at price $p_L = 100$ if the low state does occur and also that he can make at least one profitable transaction at a price $p_H = 200$ if the high state does occur. Assume further, that subjects "arrive at" the market randomly; this assumption simply entails that the probability that a person is quick at typing on the computer terminal is uncorrelated with his or her identity as a buyer or a seller.

Recall that each trading period in the double auction lasts for four minutes. After a few seconds, a buyer and seller will have "arrived at" the market, and there will be a posted bid for 100 and an offer 200 since these strategies entail no risk. *Assume now that the subjects attach no informational value to the sunspot announcement*. Look again at Table 4, and notice that there are two buyers (B4 and B5) and two sellers (S4 and S5) who actually prefer the high state to the low state. Likewise, there are two buyers (B1 and B2) and two sellers (S1 and S2) who prefer the opposite.[13]

The notion that the "stalemate" in a computerized double auction--with a standing bid at 100 and a standing offer at 200--will last for four minutes is ludicrous. Indeed, it is not even individually rational for any agent who can induce an information cascade that reinforces his own private forecast. Thus it is perfectly reasonable to assume that some buyer (either B4 or B5) will accept an offer at 200 or that some seller (either S1 or S2) will accept a bid at 100. In either case, a transaction will occur, and it doesn't matter which one actually does happen.

[13] Indeed, only B3 and S3 are indifferent about the states, as long as each transacts at the *right* price for the *right* state.

This first transaction can now be modeled easily as the beginning of an information cascade. Let the expression $prob\{H \mid p_1 = 100\}$ denote the probability that the median transaction will be high conditional on the first transaction having occurred at a price of 100. Under the (very) conservative assumption that each subject thinks that future high and low transactions are still equally likely, this event occurs only if at least three of the remaining five transactions occur at $p_H = 200$.[14] Hence,

$$prob\{H \mid p_1 = 100\} = \sum_{k=4}^{5} \binom{5}{k} (1/2)^k (1/2)^{5-k} + (1/2)\binom{5}{3}(1/2)^3(1/2)^2 = 11/32.$$

The first term on the right is the probability that there are four or five transactions at a high price. The last term in this expression assigns probability one-half to the event that the state is high if there have been three low transactions and three high ones. Likewise, if $p_H = 200$ has been the first transaction price to occur, then the probability that the median transaction will be low is $prob\{L \mid p_1 = 200\} = 11/32$.

Assume now that a bid of $p_L = 100$ and an ask of $p_H = 200$ again appear on the subjects' screens. Now the $i-th$ agent—buyer or seller--who agrees to a price of $a_i \in \{100,200\}$ will expect to earn $\psi(p_1, p_2)v_i(H, a_i) + (1-\psi(p_1, p_2))v_i(L, a_i)$, where $\psi(p_1, p_2) = prob\{H \mid p_1, p_2\}$ and now $v_i(s, a_i)$ is the rent earned if he takes action $a_i \in A_i$ and state $s \in \{H, L\}$ occurs. Using the data in Table 2, it is easy to check that if $\psi(p_1, p_2) \leq 0.21$, then no seller has negative expected profits from making an offer at 100. Analogously, if $\psi(p_1, p_2) \geq 0.79$, then no buyer has negative expected profit from purchasing at 200. *Our simple behavioral assumption is that any agent will make a*

[14] Even though the treatments all specify that a median of 150 implies that high has occurred, it is more realistic in this highly discrete action space to assume that high or low states occur with equal probability if

transaction as long as the expected profits are positive. This assumption simply states that agents prefer to make some money than to continue a stalemate in which bids and offers differ widely and never the twain shall meet.

The key step in the cascade now follows. After one transaction has occurred, the next transaction at the same price pushes this probability into the critical range $\psi(p_1, p_2) \in [0,0.21] \cup [0.79,1]$. Since

$$prob\{H \mid p_1 = 100, p_2 = 100\} = \sum_{k=4}^{4} \binom{4}{k} (1/2)^k (1/2)^{4-k} + (1/2)\binom{4}{3}(1/2)^3 (1/2)^1 = 3/16,$$

a second transaction at a low price will insure that every agent has an incentive to bid and offer at $p_L = 100$ because now no one's expected profits are negative. Likewise, $prob\{H \mid p_1 = 200, p_2 = 200\} = 13/16$ if two high transactions have occurred. Imposing the principle of backward induction, we see that there are two equilibria in this simple model of an informational cascade: if the initial transaction occurs at a low price, then there is an equilibrium in which all transactions are at $p_L = 100$, and if the initial transaction occurs at a high price, then there is another equilibrium in which the next five transactions all occur at $p_H = 200$. Finally, note that our behavioral assumption precludes anyone from making a second transaction that "leans against the wind" because to do so would entail one side—either the buyer or the seller--having negative expected profits.[15] So the information cascade is robust, once the first random transaction has occurred.

---

exactly three transactions at 100 and three at 200 take place.

[15] For example, if the first transaction has occurred at a low price, then a second transaction that leans against the wind would of necessity have a buyer purchasing at a high price. After two offsetting transactions, the posterior probability of the state being high would then be 50%. But the rents in Table 4 show that *every* buyer has negative expected profits from buying at 200 if his posterior probabilities on the state being high are only 50%.

Notice how little we have assumed about the information that the subjects use. They do not need to know all the valuations in the treatment; all they do need to know is that there are six transactions in a period. Transaction volume information is common knowledge in the MUDA software, and subjects can draw on their experience in the six initial rounds, in which volume is almost always equal to 6 units. Further, we do not need to make strong rationality assumptions. All that is necessary is that several (not all) of the subjects have a sense of backward induction and that everyone prefers transacting at a price where expected profits are positive over maintaining a stalemate where no transaction occurs!

6.2 Corroborating Evidence

We examined the data from the four double auction sessions where a sunspot equilibrium did not obtain. We focused on the very first transaction in periods where the sunspot announcement was not self-fulfilling. Table 5 summarizes our findings.

Table 5: Data on the First Transaction When the Sunspot is Wrong

Double Auction Session	Sunspot Differs from Realized State	First Transaction *is Not* Consistent with Realized State (Period, Buyer-Seller, Price)	First Transaction *is* Consistent with Realized State (Period, Buyer-Seller, Price)
Experimenter Announcements Session 1	4 times	Period 7, B5-S1, 145	Period 12, B1-S2, 149 Period 14, B3-S1, 130 Period 15, B2-S5, 110
Experimenter Announcements Session 2	5 times	Never	Period 8, B1-S2, 130 Period 11, B3-S2, 120 Period 12, B3-S2, 110 Period 14, B3-S4, 110 Period 15, B3-S2, 109
Experimenter Announcements Session 3	2 times	Period 10, B1-S4, 110	Period 12, B2-S3, 120
Coin-Flip Announcements Session 3	5 times	Never	Period 7, B3-S3, 120 Period 8, B5-S1, 102 Period 9, B4-S1, 111 Period 11, B5-S1, 101 Period 15, B4-S2, 105
Totals	16 periods	2 periods	14 periods

We see that in 14 of the 16 trading periods in which the sunspot announcement did not correspond to the realized state, the first transaction price *was* consistent with the realized state. Indeed, a binomial test based upon the null hypothesis that the first transaction price has no predictive power has a p-value of 0.02. So we can (easily) reject this null hypothesis for a test of size 1%. We note further that the initial transaction price "sets the tone" for all further transaction prices within each of these trading periods; in all 14 cases, the subsequent transaction prices were all consistent with the state of the world signaled by the first transaction price.

The data in Table 5 also allow us to examine whether, sellers S1 and S2 are more likely to transact at an initially low price or whether buyers B4 and B5 are more likely to transact at an initially high price. Recall from Table 4 that sellers S1 and S2 earn the highest profits if the low state obtains while buyers B4 and B5 earn the highest profits if the high state obtains. Hence, these subjects have the greatest incentives to make the first transaction and start an information cascade. From Table 5 we see that in all 14 cases where the sunspot was not self-fulfilling and the first transaction price correctly signaled the realized state, this first transaction price was always in the domain of the low state. Thus in these 14 cases the sunspot announcement was high, but the realized state turned out to be low. Confirming our model of an information cascade, we see that in 10 of these 14 cases, the seller involved in the first transaction was either S1 or S2, the two who had the greatest stake in ensuring that a low state of the world was realized.

We note further that in the two cases where the sunspot was not self-fulfilling and the first transaction did not correctly signal the realized state, the second transaction *did* correctly signal the realized state. In period 7 of the Double Auction, Experimenter Announcements, Session 1, the sunspot announcement was low, but the state turned out to be high. As noted in Table 5, the first transaction was at a price of 145, in the domain of the low state, and involved buyer B5. This was followed by a second transaction, also involving buyer B5, at a price of 155, which lies in the domain of the high state. The subsequent transaction prices were 175, 180, 175, 192, 185. Thus, consistent with our theory, buyer B5 was instrumental in starting a contagion towards the high state realization, but it began with the second transacted price. Similarly, in period 10 of Double Auction, Experimenter Announcements, Session 3, the sunspot announcement was low, but the state turned out to be high. As Table 5 reveals, the first transaction was at a price of 110, in the domain of the low state. However, the second transaction at a price of 210 correctly signaled the resulting high state and involved buyer B5. The third transaction was at a price of 165 and involved buyer B4. Again, this behavior is strongly consistent with the predictions of our theory.

Finally, we note that in 42 of the 44 double-auction trading periods in which the sunspot announcement *was* self-fulfilling, the first transaction price was always consistent with the realized state. That is, the first two players to complete a transaction acted as though they believed the sunspot announcement would be self-fulfilling.

7. Sunspots are Correlated Equilibria in the Call Markets

Cass and Shell (1983) and Azariadis (1981) originally developed the theory of sunspots in a competitive general equilibrium setting with infinitely many agents. By contrast, the market environments we study necessarily involve a finite number of subjects. As we have seen, one player in a double auction can precipitate an information cascade that affects other players' beliefs about the likely state. The purpose of this section is to show that such behavior is *not* possible in the call market we examine.

In the standard call market, players know only their own private values or costs and these are not state-dependent. Then the individual's problem is to choose a bid or an offer taking into account that this action may affect both the price and the probability that the individual is able to buy or sell. Since the reservation values of the other subjects are not common knowledge, the game is one of incomplete information, and Bayesian Nash equilibrium is the appropriate solution concept. Such a model is beyond the scope of this paper, but we do draw upon such theoretical analyses surveyed in Satterthwaite and Williams (1993). An important theoretical issue is the extent to which a buyer's shade their bids and sellers boost their offers. Restricting their attention to a homogeneous and also to units that actually have a chance of being traded, Rustichini, Satterthwaite and Williams (1994) show that this kind misrepresentation would be no larger than 20% in our experiments if the state were common knowledge

It seems plausible then to impose that each buyer has an action set $A_i = \{v_i(H), v_i(L)\}$, where $v_i(s) = (v_{i,1}(s), v_{i,2}(s))'$ is the 2×1 vector of induced values in state $s \in \{H, L\}$. Likewise, each seller has an action set $A_i = \{c_i(H), c_i(L)\}$, where now $c_i(s) = (c_{i,1}(s), c_{i,2}(s))'$ is a 2×1 vector of state-contingent costs. Let N be the set of players. Then we can follow the usual notional convention and write the profile of actions $a = (a_i, a_{-i}) \in A = \times_{i\in N} A_i$. Then each player's payoff is:

$$u_i(a_i, a_{-i}) = \begin{cases} \sum_{j=1}^{2} \chi_{i,j}(a_i, a_{-i})[v_{i,j}(a_i, a_{-i}) - p(a_i, a_{-i})] & \text{if } i \text{ is a buyer} \\ \sum_{j=1}^{2} \chi_{i,j}(a_i, a_{-i})[p(a_i, a_{-i}) - c_{i,j}(a_i, a_{-i})] & \text{if } i \text{ is a seller} \end{cases}$$

where $\chi_{i,j}(a_i, a_{-i})$ is an indicator function that takes on a value of 1 if the i-th agent trades his j-th unit and 0 otherwise. *It is important to notice that the state of nature does not enter into these payoffs, since the agents' choices completely determine it.* A *correlated equilibrium* of the strategic game $\langle N, (A_i), (u_i)\rangle$ is a finite probability space (Ω, π), an information partition $\Im_i$ for every $i \in N$, and strategy $\sigma_i : \Omega \to A_i$ for every $i \in N$ that satisfy the two properties. First, $\sigma_i(\omega) = \sigma_i(\omega')$ whenever $\omega \in P_i$ and $\omega' \in P_i$ for some $P_i \in \Im_i$. Second, for every $\tau_i : \Omega \to A_i$ such that $\tau_i(\omega) = \tau_i(\omega')$ whenever $\omega \in P_i$ and $\omega' \in P_i$ for some $P_i \in \Im_i$, the following inequality holds:

$$\sum_{\omega\in\Omega} \pi(\omega) u_i(\sigma_i(\omega), \sigma_{-i}(\omega)) \geq \sum_{\omega\in\Omega} \pi(\omega) u_i(\tau_i(\omega), \sigma_{-i}(\omega)),$$

where again the notation $\sigma_{-i}(\omega)$ follows the usual convention.

7.1 A Correlated Equilibrium

We will now show that call market data we observed were quite close to a correlated equilibrium. First, set $\Omega = \{H, L\}$ with the interpretation that a state of nature is just the random announcement. Of course, the probability measure has $\pi(\{H\}) = \pi(\{L\}) = 1/2$ since it was common knowledge that either announcement was equally probable. Second, let each agent's information partition be $\Im_i = \{\{H\}, \{L\}\}$ for all $i \in N$ since the announcement was public in both treatments. Third, each agent adopts the following strategy

$$\sigma_i(\omega) = \begin{cases} v_i(\omega) \text{ if } i \text{ is a buyer} \\ c_i(\omega) \text{ if } i \text{ is a seller} \end{cases}$$

where again $\omega \in \Omega$ is the random realization of the announcement.

We will now analyze every possible unilateral deviation by any agent. Look closely at the induced values and costs illustrated in Figure 1, and recall that the price in a call market is $\tilde{p} = (\underline{p} + \overline{p})/2$. Consider first the state $\omega = L$. If the announced state is low, then the rule for price determination is such that no unilateral deviation can change the state, but a deviation does have a marginal effect on $\tilde{p}$ and thus every player's payoffs. Every buyer who deviates by declaring his high values will lose profits; indeed buyers B1 through B4 must purchase a second unit at a loss, and buyer B5 still buys two units, but she raises the market price. Every seller who deviates by declaring high costs will sell nothing and thus give up the rents she would otherwise have earned. Consider second the state $\omega = H$. Any buyer who deviates by declaring low values will buy nothing and give up all the rents she would have otherwise earned. If any seller S2 through S5 deviates by declaring low costs, she will be forced to sell a second unit at a loss. And if seller S1 deviates analogously, she will still sell both units but at a lower price, decreasing the rent she would have otherwise earned. We conclude these strategies constitute a *correlated equilibrium* according to Aumann's (1974) definition.

The same need not be true in the double auction, as we have shown with our model of an information cascade in Section 6. The main difference is that in a double auction bids, asks, and transaction prices are revealed as they occur during a trading period. In the call market, submitted bids and asks are sealed, and price determination is centralized, an aspect of the environment that rules out the possibility of an information cascade being started by a single trader.

7.2 Corroborating Evidence

We examined our data from the call market sessions to see whether the assumptions made in the preceding section are reasonable. On the one hand, the fact that the sunspot announcement always predicted the resulting state in every round of every call market session should be sufficient evidence to corroborate our story. On the other hand, it is important to check how reasonable our assumption that there is not much bid shading or offer boosting.

Let $\sigma_{i,j}(\omega)$ denote the strategy of player *i* on unit *j* in state $\omega \in \Omega = \{H, L\}$. Then the percentage misrepresentation for a buyer is $[v_{i,j}(\omega) - \sigma_{i,j}(\omega)]/v_{i,j}(\omega)$, and that for a seller is $[\sigma_{i,j}(\omega) - c_{i,j}(\omega)]/c_{i,j}(\omega)$. Table 6 reports the median (and the averages) of these misrepresentation percentages over the last ten trading periods—the ones where agents faced uncertainty about the state of nature--for every call market session.

Table 6: Median Percentage (and Average) Misrepresentation

Session	Buyers		Sellers	
	First Unit	Second Unit	First Unit	Second Unit
Experimenter Announcements #1	1.6% (-1.9%)	0.3% (3.6%)	1.4% (10.1%)	1.1% (9.4%)
Experimenter Announcements #2	4.6% (-0.6%)	0.5% (7.0%)	2.6% (22.9%)	1.1% (8.6%)
Experimenter Announcements #3	4.2% (2.5%)	6.7% (10.7%)	2.0% (15.3%)	0.6% (22.3%)
Coin Flip Announcements #1	3.2% (-11.4%)	1.4% (12.1%)	17.9% (27.2%)	0.8% (4.2%)
Coin Flip Announcements #2	3.1% (1.1%)	3.1% (6.5%)	8.1% (13.5%)	0.9% (11.1%)
Coin Flip Announcements #3	8.9% (12.2%)	6.3% (22.2%)	12.1% (18.5%)	3.0% (11.5%)

The median percentage misrepresentation amounts are all positive, and they are all well within the theory developed by Rustichini, Satterthwaite and Williams (1994) for a simpler model with homogenous buyers and sellers. The averages in each session are slightly greater in absolute value, indicating that some agents systematically shade bids and boost offers. But even these averages are not so egregious as to cast serious doubt on the model's assumption that the agents report their values or costs truthfully in the correlated equilibrium. Also, for completeness, we ought to mention that the medians and averages for the second units don't matter as much since only one buyer and one seller in any state of nature actually trades two units in equilibrium. We conclude that there is strong evidence that data from these call markets seem to indicate that the agents are following a correlated equilibrium. We believe this is the first empirical verification of the notion that a sunspot equilibrium is indeed a correlated equilibrium of a market game.

8. Conclusions

Experimental economics has been perhaps most successful in describing now markets work. It has shown unequivocally that institutions matter: different kinds of markets give rise to different outcomes. Until now, models whose equilibria relied upon the existence of "animal spirits" have been useful and elegant theoretical curiosities. But we have given these models real empirical bite. We have shown that extrinsic uncertainty is an important part of real markets. It was surprising to us that the flow of information is an important determinant of whether sunspot equilibria regularly obtain and this finding has obviously important theoretical implications. A sunspot announcement is a reliable coordinating device when information flows slowly, as in a closed book call market. The theory of sunspot equilibrium has been developed typically in a Walrasian framework (as in Cass and Shell (1983)) or in a dynamic framework (as in Azariadis (1981)), but in either case the flow of information was no slower nor faster than the speed at which a market cleared. It seems now that it is important to model how information flows while the market is clearing.

Our experimental design also suggests that it is important to consider the semantics of the language of sunspots. When we had subjects flip a coin, we asked them to state that the forecast is high if heads came up. This is quite different from stating that heads has come up and then expecting the subjects to coordinate by themselves, as though every person can come up (simultaneously) with a semantic interpretation of what the event heads might mean. Thus it may not be enough to train subjects with a blinking light on the screen, but it certainly is adequate to state, "The forecast is high." Likewise, there may be a sunspot in the NYSE based upon whether the NFC or the AFC wins the Super Bowl, but there will be no empirically testable hypothesis until it has become common knowledge that the language of that particular sunspot depends upon everyone knowing which teams belong to each conference.

The fact that sunspot announcements serve as coordinating devices has important implications for financial markets in the field. Indeed, one might argue that an important aspect of monetary policy in the United States in the last few years has been trying to anticipate and perhaps mitigate the effects of sunspot equilibria in major financial markets. Thus showing that sunspot equilibria may well depend on the flow of information has very real implications for the architecture of these markets. For example, our findings would seem to indicate that stock trade suspension rules that are commonly used centralized markets in the field actually increase the possibility of sunspot equilibria.

In these experiments, the sunspot realizations help, paradoxically, to ensure that the market is fully efficient. Indeed, when agents try unsuccessfully to manipulate information strategically, they pay a price because the wrong infra-marginal bids or offers are submitted. This implies that that the resultant equilibrium does not maximize social surplus, conditional upon the state of nature that obtains. In field markets, the connection between welfare and sunspots is not so apparent, but it may be a very real part of any financial system.

References

Aumann, Robert .J. "Subjectivity and Correlation in Randomized Strategies." *Journal of Mathematical Economics* 1 (1974), 67--96.

Azariadis, Costas. "Self-Fulfilling Prophecies." *Journal of Economic Theory* 25 (1981), 380-96.

Ball, Sheryl, Catherine Eckel, Philip J. Grossman and William Zame, "Status in Markets," forthcoming in *Quarterly Journal of Economics,* 2001.

Benhabib, Jess and Roger E.A. Farmer. "Indeterminacy and Sunspots in Macroeconomics," in J.B. Taylor and M. Woodford (eds.) *Handbook of Macroeconomics* v. 1A, New York: North-Holland, 1999, 387-448.

Brandts, Jordi and Charles Holt. "An Experimental Test of Equilibrium Dominance in Signaling Games," *American Economic Review* 82 (1992), 1350-1365.

Brandts, Jordi and W. Bentley MacLeod. "Equilibrium Selection in Experimental Games with Recommended Play," *Games and Economic Behavior* 11 (1995), 36-63.

Cass, David and Karl Shell. "Do Sunspots Matter?" *Journal of Political Economy* 91(1983), 193-227.

Farmer, Roger E. A. *The Macroeconomics of Self-Fulfilling Prophecie*s. Cambridge, MA: The MIT Press, 2 nd Edition, 1999.

Flood, Robert P. and Peter M. Garber. "Market Fundamentals versus Price-Level Bubble: The First Tests," *Journal of Political Economy* 88 (1980), 745-70.

Jevons, William S. Investigations in Currency and Finance. London: Macmillan, 1884.

Mackay, Charles. Extraordinary Popular Delusions and the Madness of Crowds. London: Richard Bentley, 1841.

Marimon, Ramon, Stephen E. Spear and Shyam Sunder. "Expectationally Driven MarketVolatility: An Experimental Study," *Journal of Economic Theory* 61 (1993), 74-103.

Peck, James and Karl Shell. "Market Uncertainty: Correlated and Sunspot Equilibria in Imperfectly Competitive Economies," The Review of Economic Studies 58 (1991), 1011-1029.

Plott, Charles and Peter Gray. "The Multiple Unit Double Auction," *Journal of Economic Behavior and Organization* 13 (1990), 245-258.

Rustichini, Aldo, Mark A. Satterthwaite, and Steven R. Williams. "Convergence to Efficiency in a Simple Market with Incomplete Information," *Econometrica* 62 (1994), 1041-1063.

Satterthwaite, Mark A. and Steven R. Williams. "The Bayesian Theory of the k-Double Auction" in D. Friedman and J. Rust (eds.) *The Double Auction Market: Institutions, Theories and Evidence*, (1993) Reading, MA: Addison-Wesley, 99-123.

Siegel, S. and N.J. Castellan Jr. *Nonparametric Statistics for the Behavioral Sciences* 2nd Edition (1988), New York: McGraw Hill.

Van Huyck, John B. Ann B. Gillette and Raymond C. Battalio. "Credible Assignments in Coordination Games," *Games and Economic Behavior* 4 (1992), 606-626.

Figure 2: Double Auction, Experimenter Announcements, Session 1

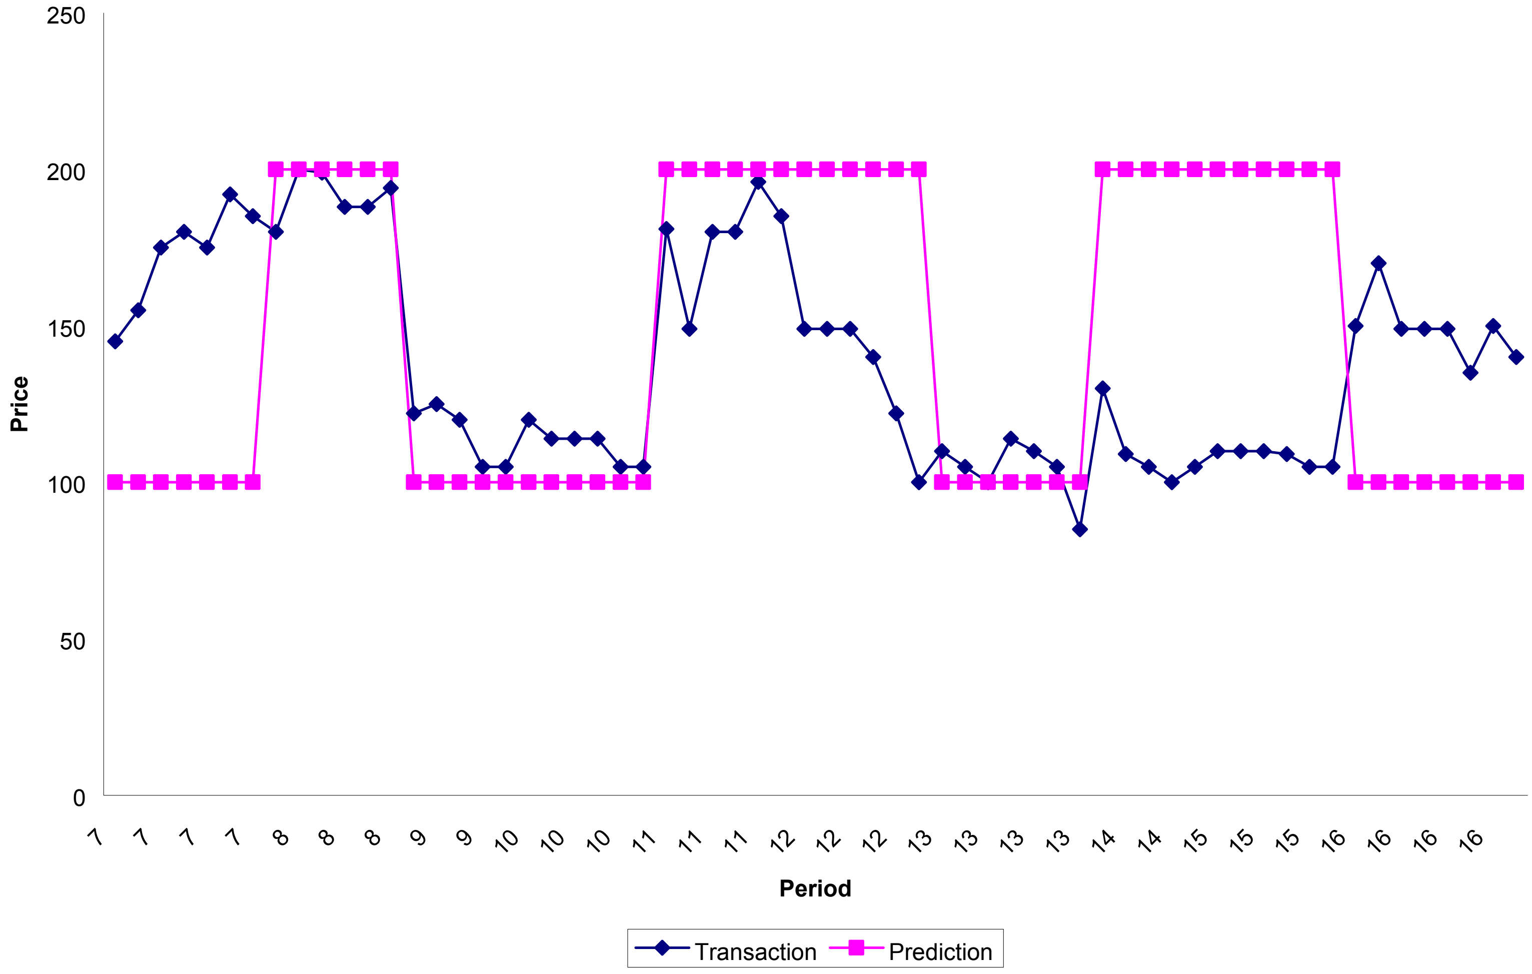

Figure 3: Double Auction, Experimenter Announcements, Session 2

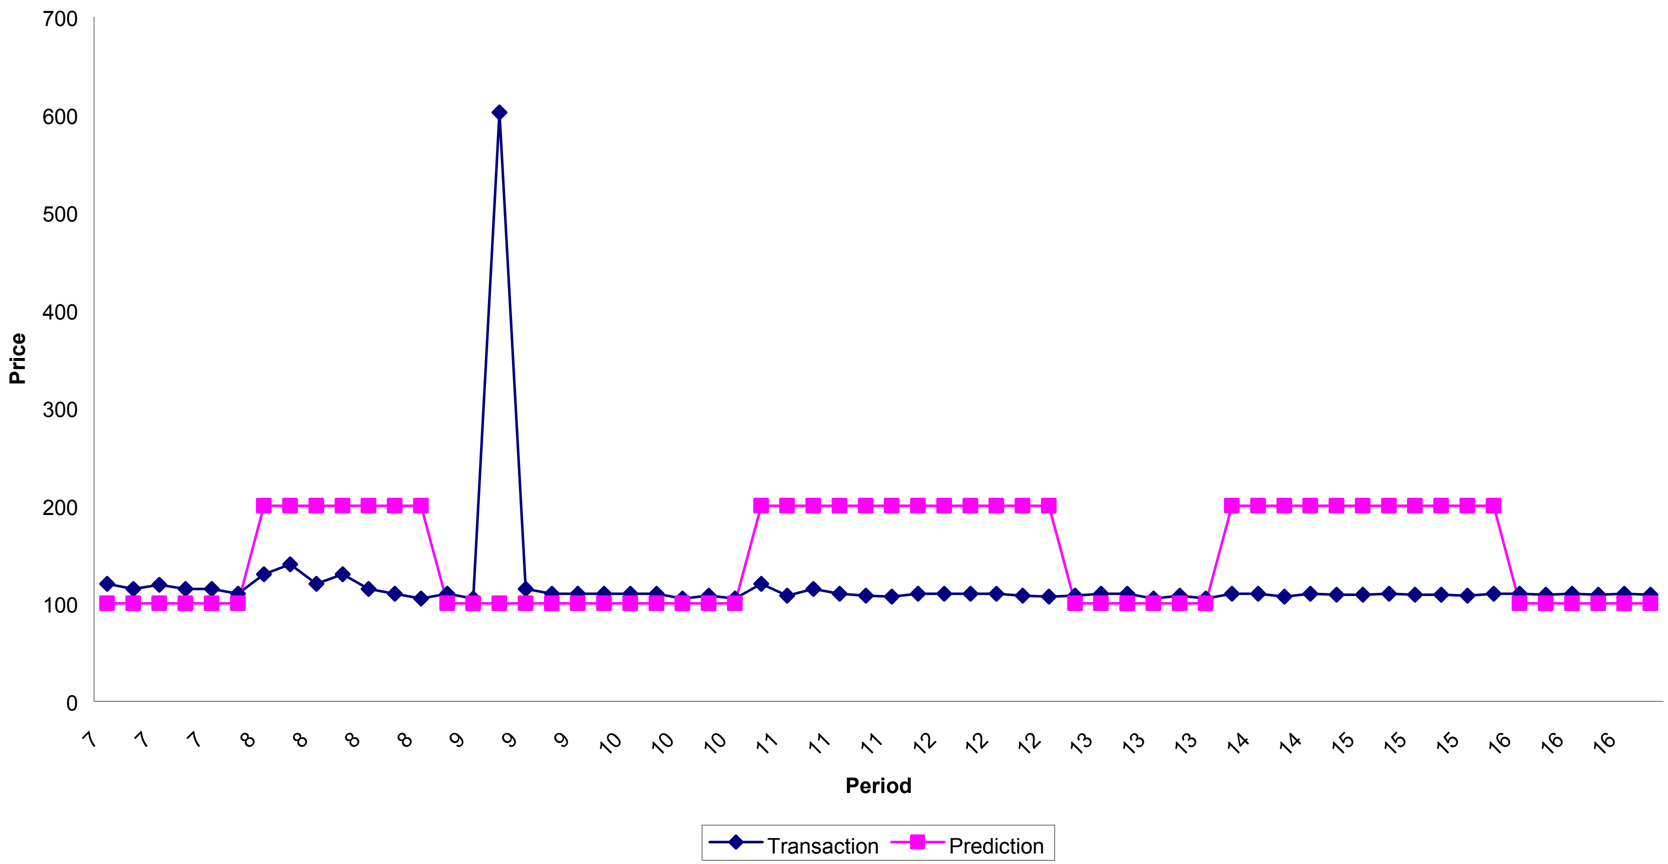

Figure 4: Double Auction, Experimenter Announcements, Session 3

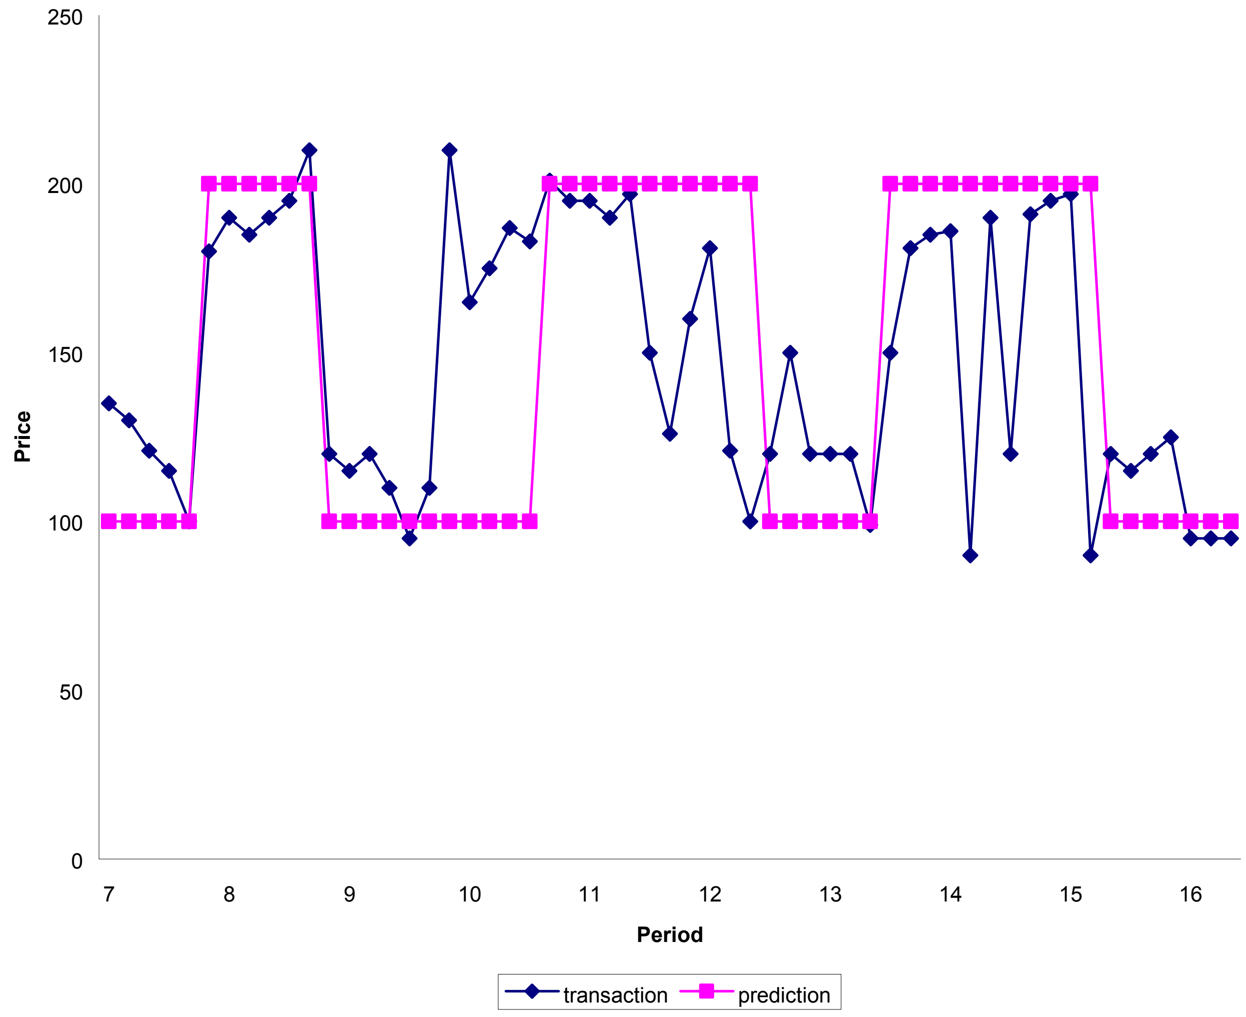

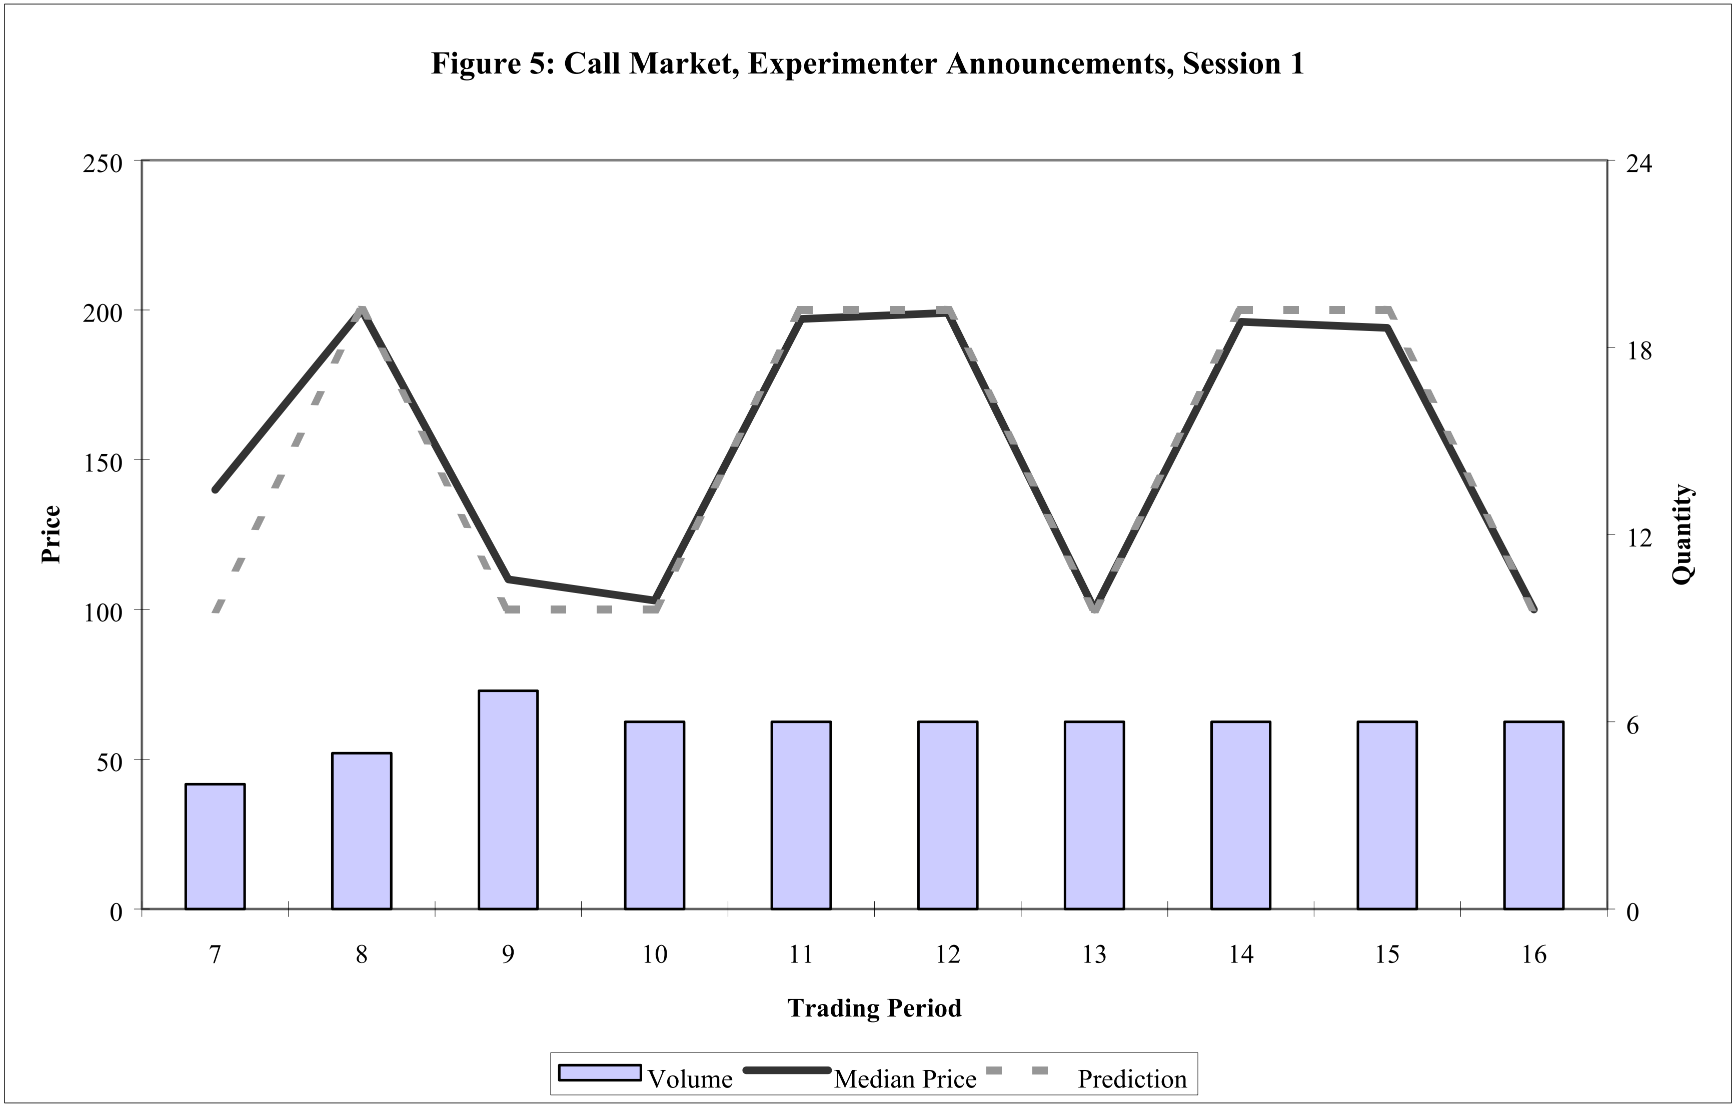

Figure 5: Call Market, Experimenter Announcements, Session 1
250
200
150
100
50
0
Price
24
18
12
6
0
Quantity
7
8
9
10
11
12
13
14
15
16
Trading Period
Volume
Median Price
Prediction

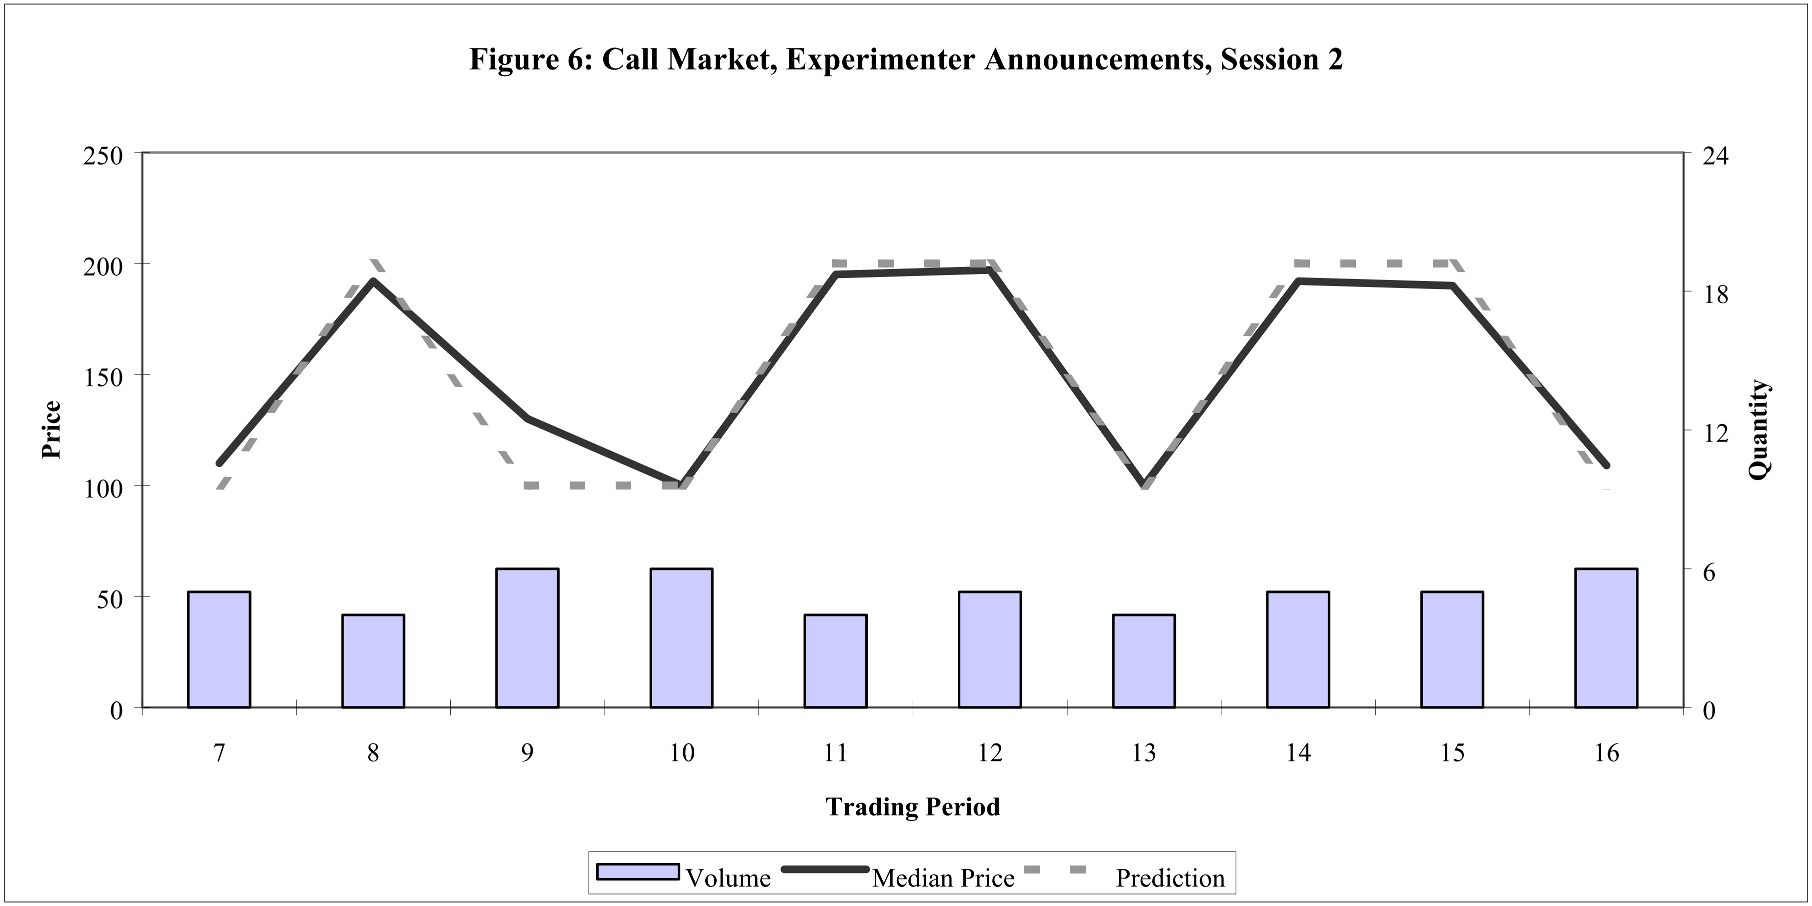

Figure 6: Call Market, Experimenter Announcements, Session 2
250
200
150
100
50
0
Price
24
18
12
6
0
Quantity
7
8
9
10
11
12
13
14
15
16
Trading Period
Volume
Median Price
Prediction

Figure 7: Call Market, Experimenter Announcements, Session 3

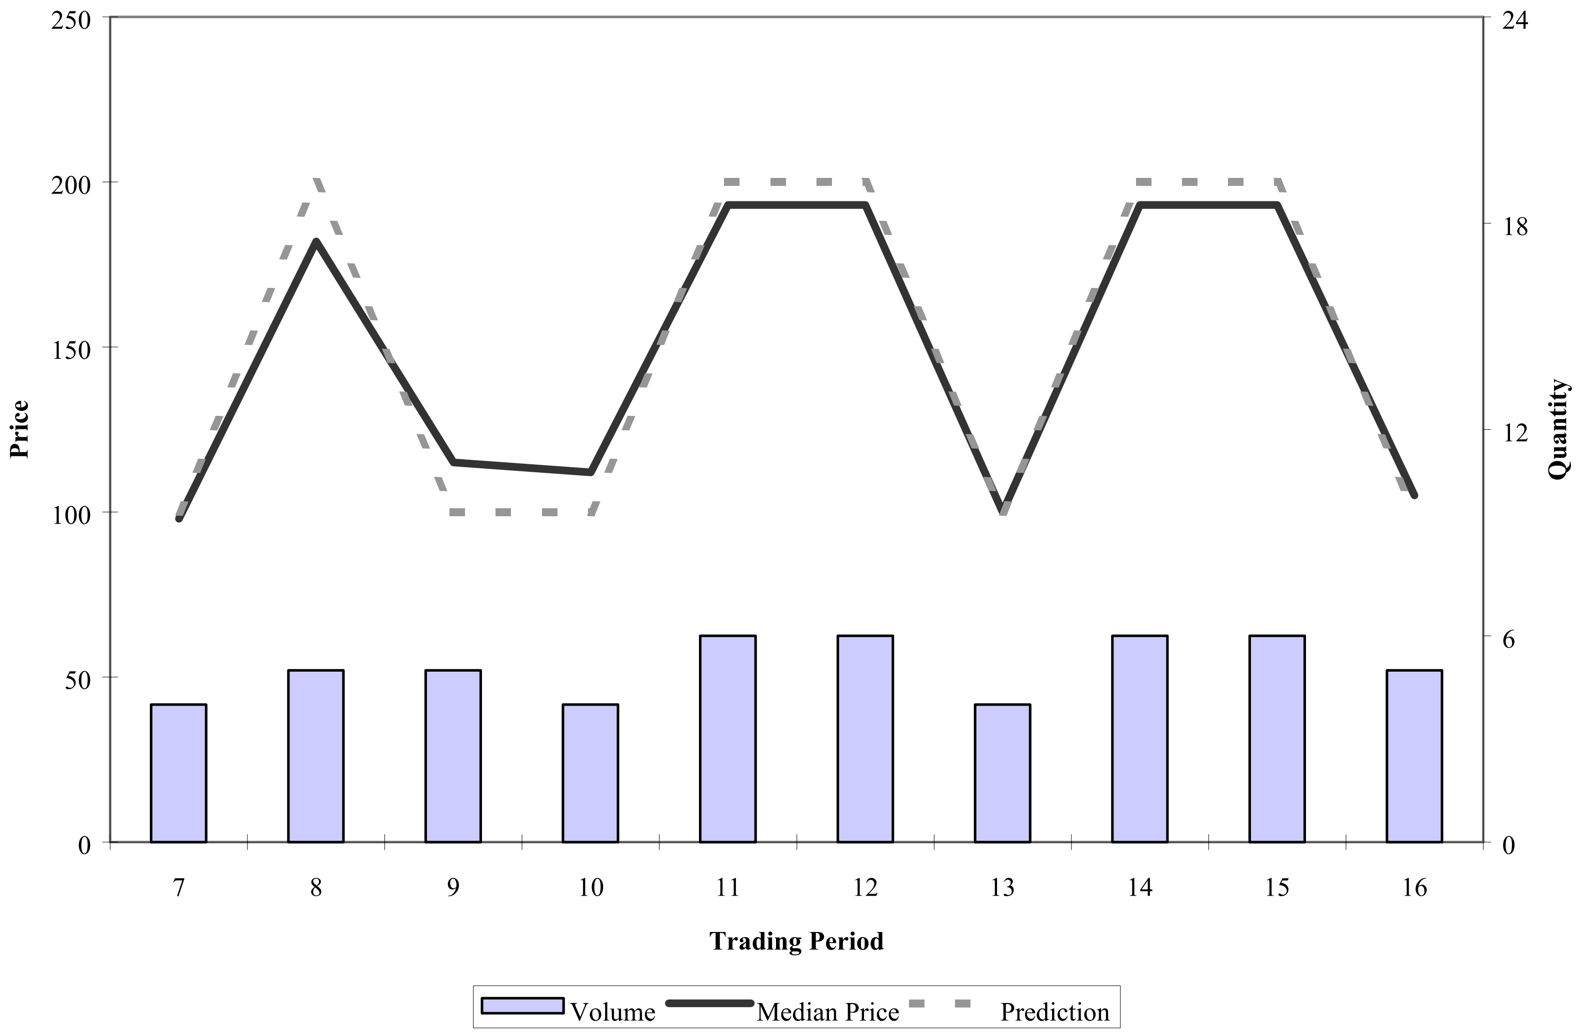

Figure 8: Double Auction, Coin-Flip Announcements, Session 1

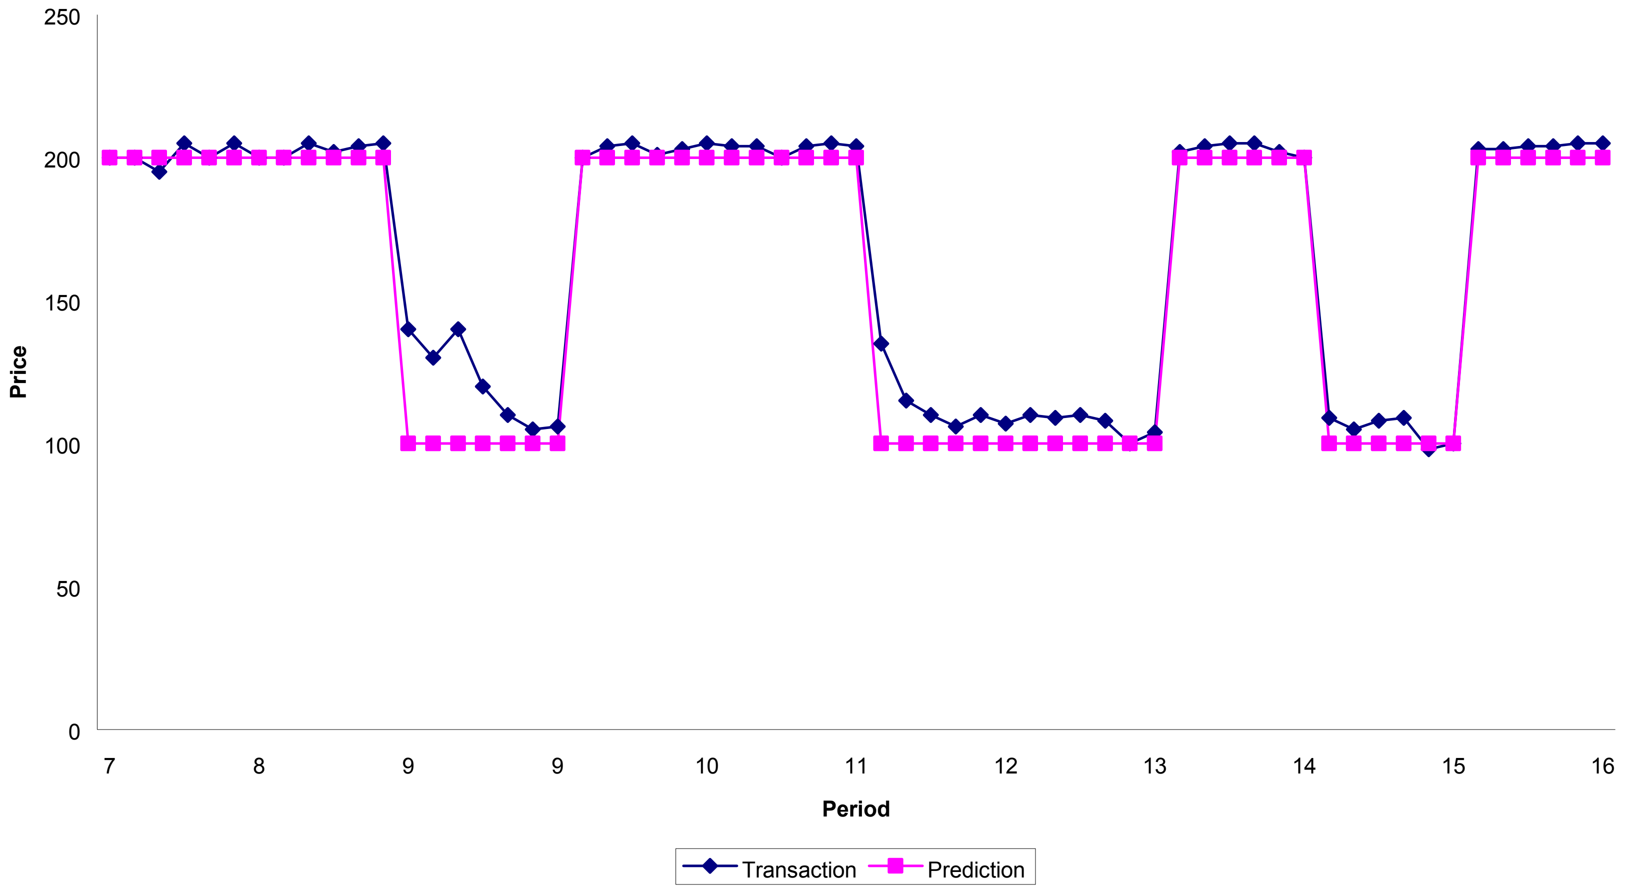

Figure 9: Double Auction, Coin-Flip Announcement, Session 2

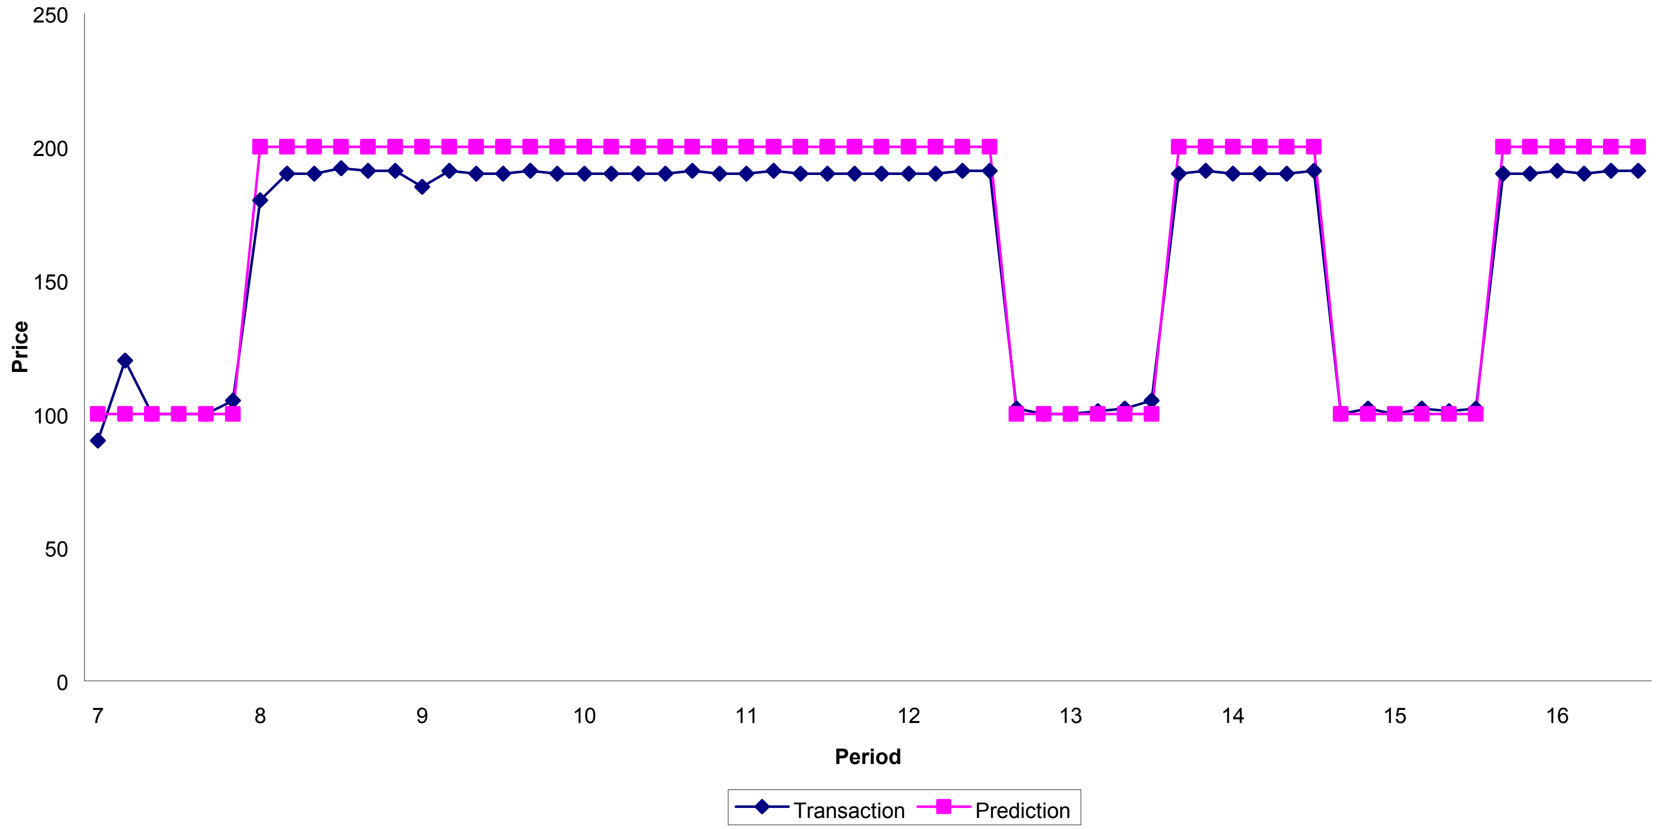

Figure 10: Double Auction, Coin-Flip Announcements, Session 3

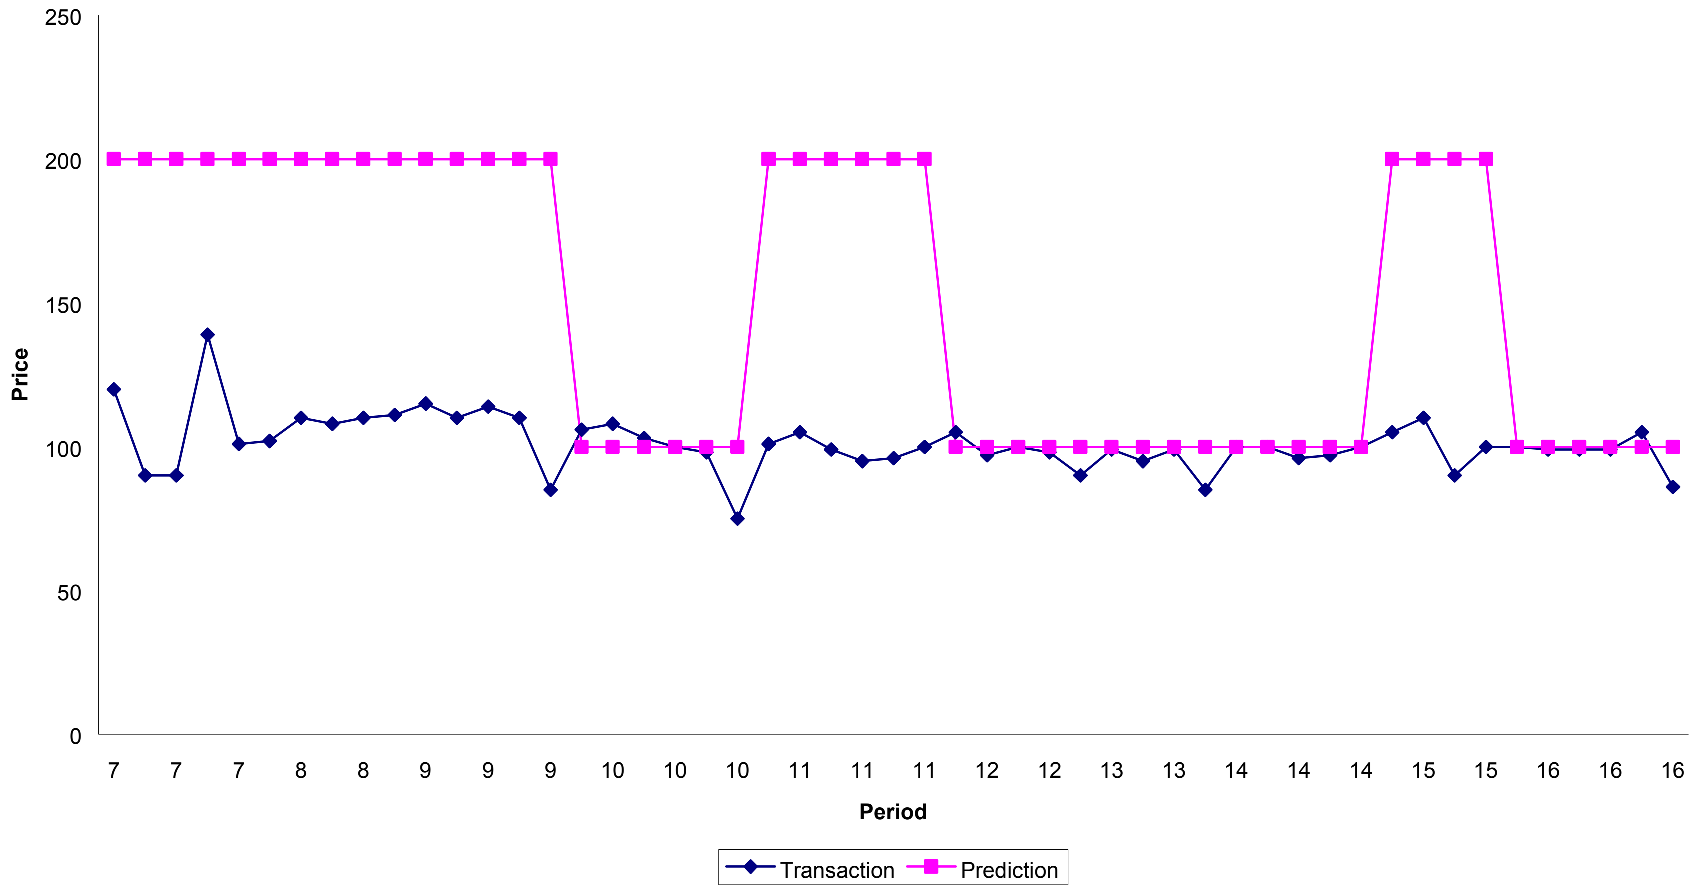

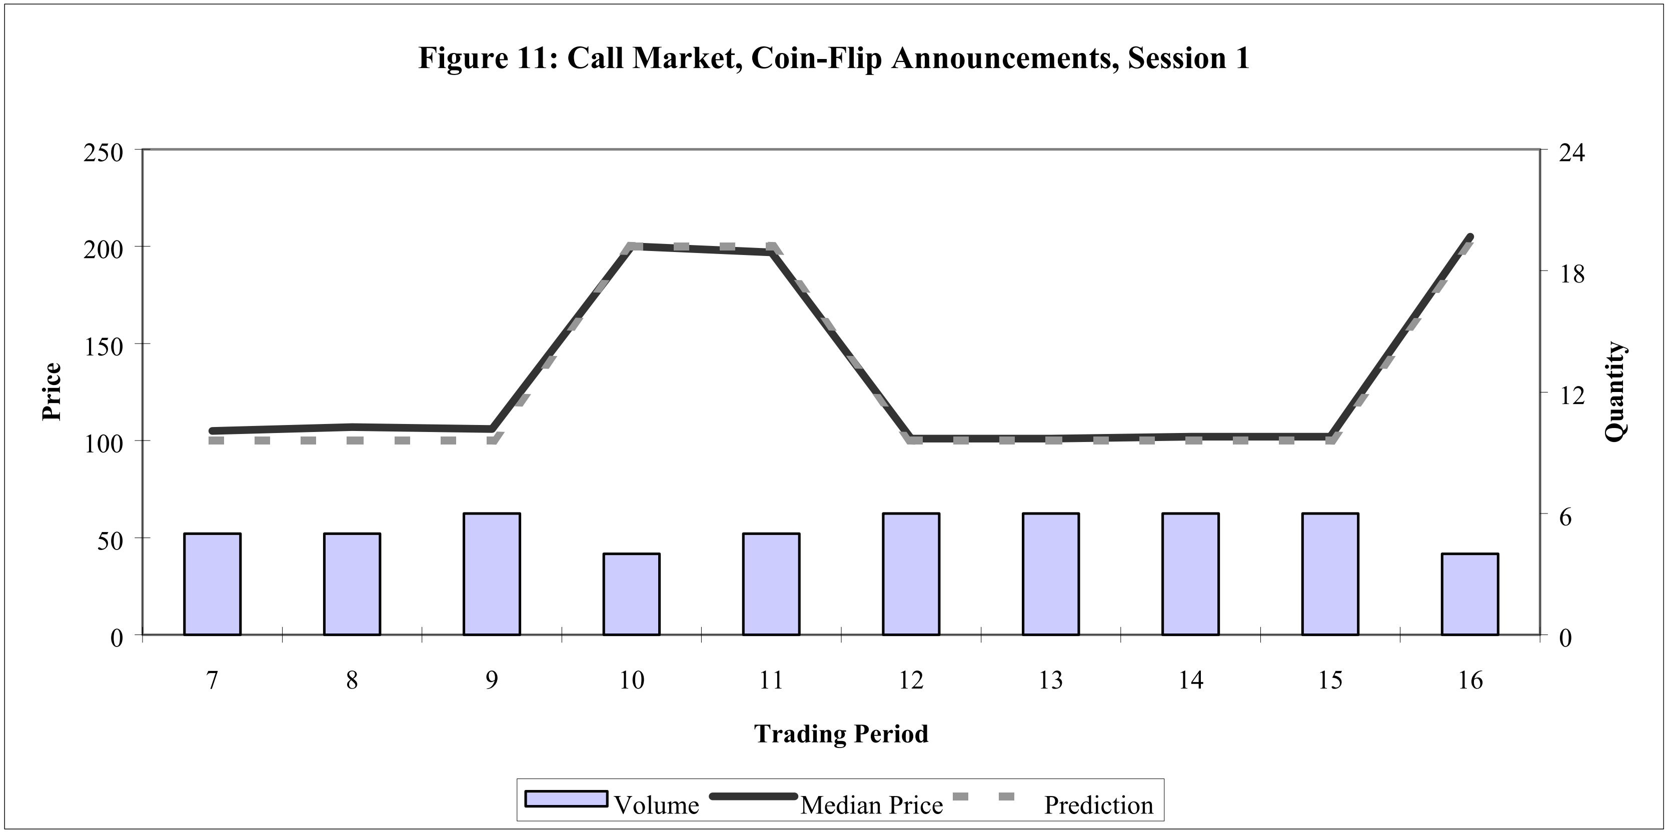
Figure 11: Call Market, Coin-Flip Announcements, Session 1
250
200
150
100
50
0
Price
24
18
12
6
0
Quantity
7
8
9
10
11
12
13
14
15
16
Trading Period
Volume
Median Price
Prediction

Figure 12: Call Market, Coin-Flip Announcements, Session 2

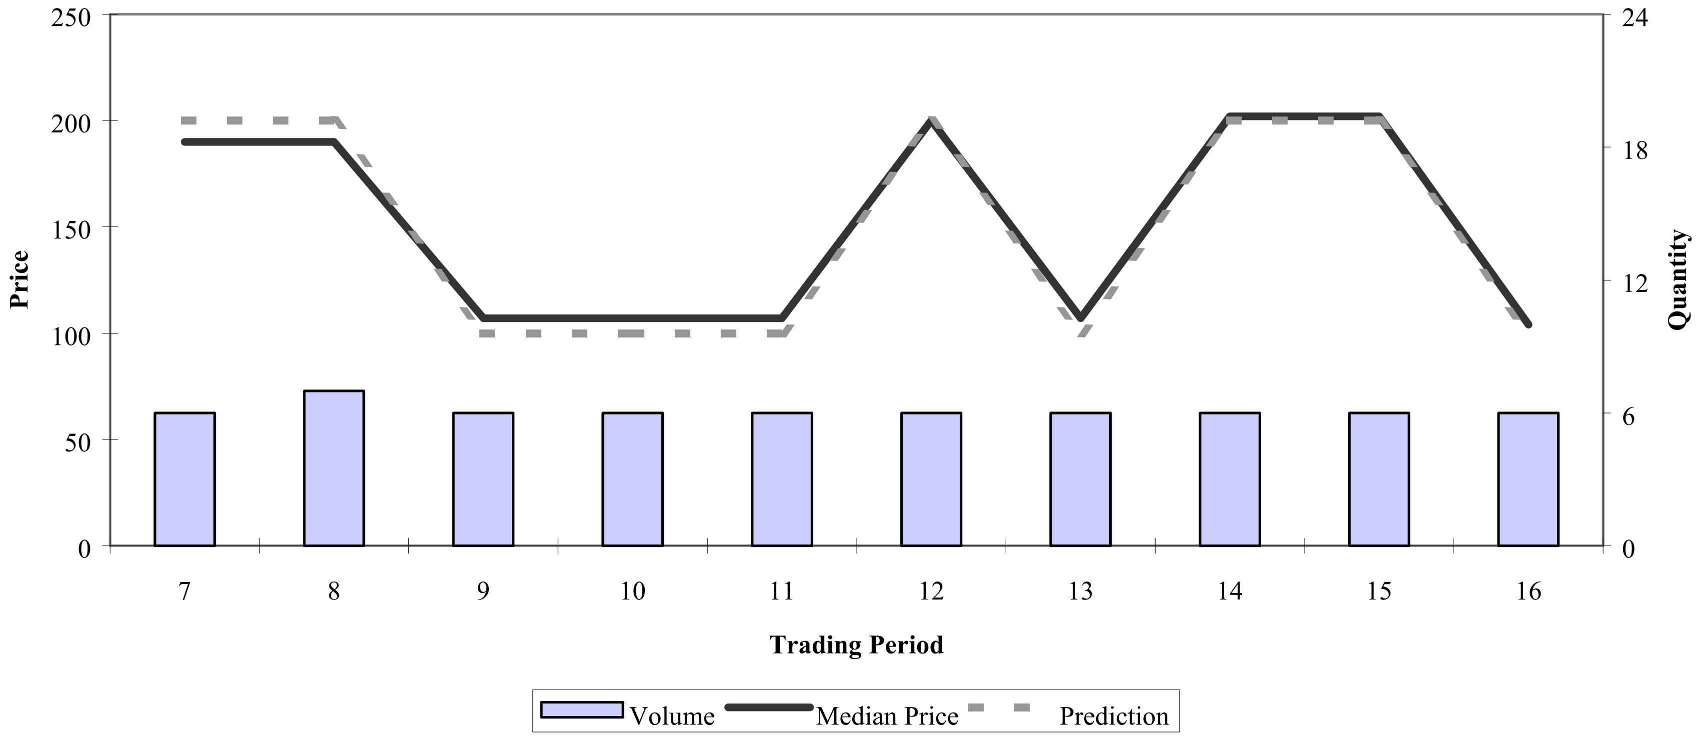

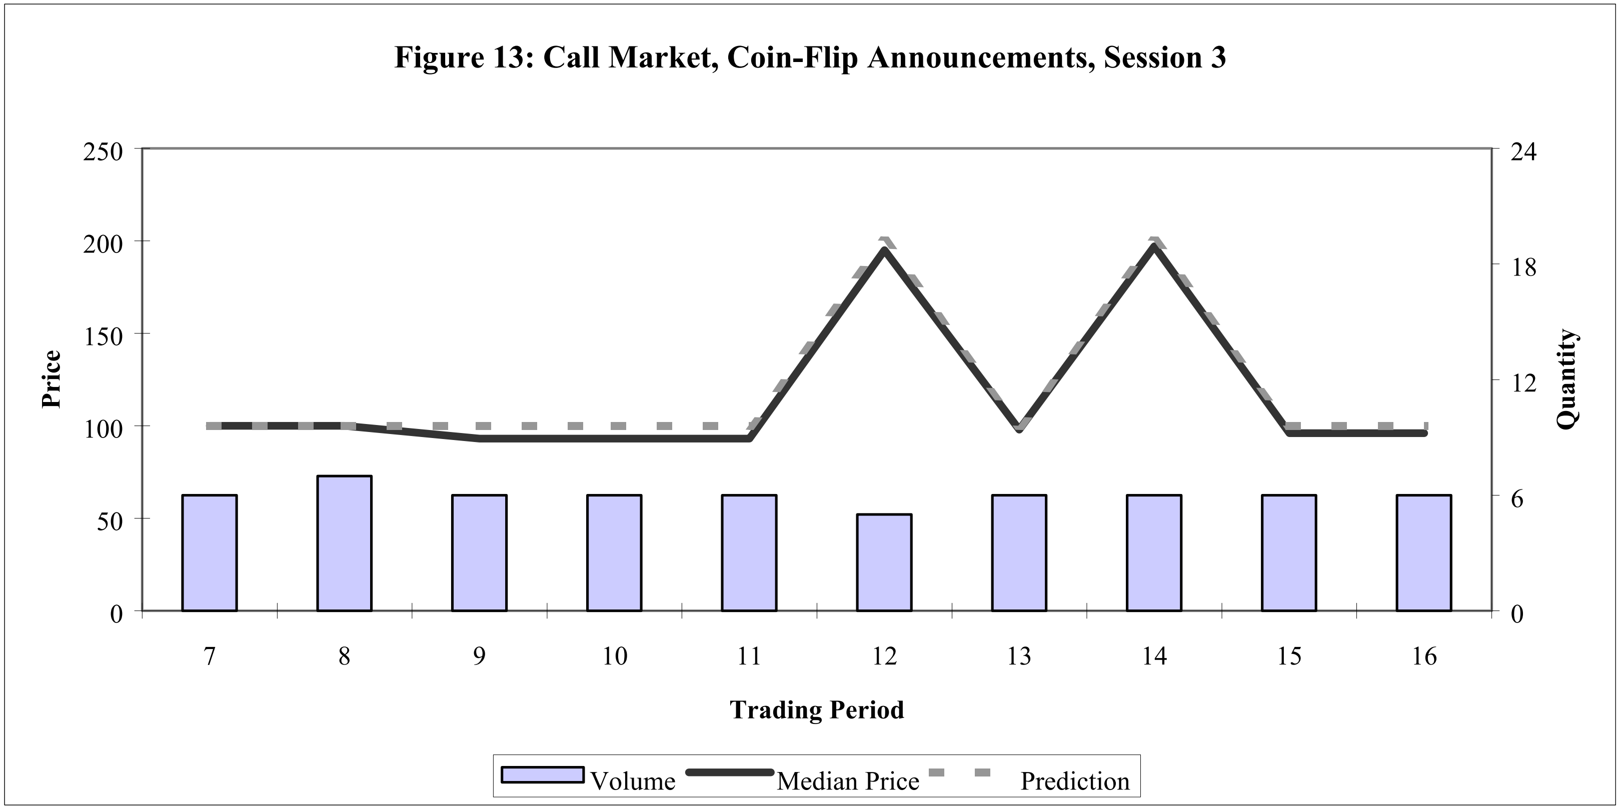
Figure 13: Call Market, Coin-Flip Announcements, Session 3
Price
250
200
150
100
50
0
Quantity
24
18
12
6
0
7
8
9
10
11
12
13
14
15
16
Trading Period
Volume
Median Price
Prediction