

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

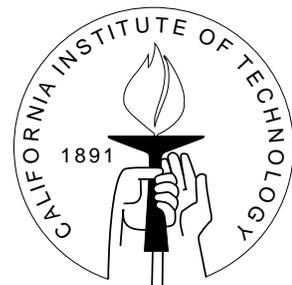
CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

PRINCIPLES OF CONTINUOUS PRICE DETERMINATION IN AN EXPERIMENTAL ENVIRONMENT WITH FLOWS OF RANDOM ARRIVALS AND DEPARTURES

Michael R. Alton

Charles R. Plott



SOCIAL SCIENCE WORKING PAPER 1256

September 2007

Principles Of Continuous Price Determination In An Experimental Environment With Flows Of Random Arrivals And Departures

Michael R. Alton

Charles R. Plott

Abstract

This paper develops a new experimental environment for the study of markets. The period structure of classical experimental markets, which is known to play an important role in the equilibration process, is replaced by an environment with a structure in which new incentives arrive randomly and continuously throughout. The issue posed is whether or not classical principles of market behavior apply in worlds where the environment is constantly changing. We develop an experimental methodology within which the issue can be studied. We report: (1) natural flow generalizations of the classical laws of demand and supply operate to dictate much of the behavior observed in the markets; (2) two different classes of laws operate: one we call the “temporal equilibrium”, which is based on the parameters that exist in the market at a moment of time, and the second is the “flow competitive equilibrium,” which reflects the underlying probabilistic structure of the parameters; and (3) the markets themselves exhibit extraordinarily high levels of efficiency.

Key words: equilibrium, experiment, price determination, continuous market, dynamic adjustment, value, complex markets

Principles of Continuous Price Determination in an Experimental Environment with Flows of Random Arrivals and Departures

Michael R. Alton

Charles R. Plott

1. INTRODUCTION¹

This paper introduces an experimental environment with complexities that have not been present in previous experimental markets. As such, the environment and procedures depart from the experimental tradition that began with Chamberlin (1948) and refined by Smith (1962). In this new environment, economic opportunities appear as a flow of incentives, on which agents can choose to act. These incentives arrive at random times, are short lived, and come from stochastic processes which change over time. In this world, supply and demand are randomly changing at each instant of time throughout the experiment.

Because such randomness and complexity appears to be a feature that is typical of environments in which markets operate, people naturally wonder about the ability of markets (and theories of markets) to cope. Fundamental questions become posed. Is the creation of such challenging environments within the capacity of existing experimental technologies? How would such experiments be conducted? Will such variability destroy the laws of markets as we know them? Do principles based on equilibrium have application or any predictive power at all in an environment that is constantly changing? Are strategic possibilities so structured that no form of equilibrium applies? Can we

¹ The authors wish to express their appreciation for the help of Roger Sherman, Kim Border, Peter Bossaerts, David Grether, and the participants in the Caltech seminar on experimental economics. The financial support of the Caltech Laboratory of Experimental Economics and Political Science and Caltech's Lee Center for Advanced Networking is gratefully acknowledged.

model the market behavior with tractable tools and in particular, is there a natural generalization of the law of supply and demand to the flow world?

The paper is divided into seven sections. The first section is this introduction. The second outlines the random arrival and departure environment that we explore. Here we develop basic ideas about the structure of the stochastic framework. The ideas turn on arrival rates of buyers and sellers, the stochastic structure of preferences of those who enter the market and the lifetime of traders while in the market. The third section is a discussion of the market institutions. The fourth section develops principles that are natural generalizations of classical principles and illustrates how they apply to the complex random arrival and departure environment. The fifth section details the experimental procedures and the experimental design that outlines the experiments that were conducted. We follow an exploratory methodology focused on broad questions, rather than a theory testing methodology. Traditionally, the exploratory approach has been employed successfully in the study of experimental markets for which theories either do not exist or lack the observational specificity required for detailed tests.² The sixth section contains the results, and the final section contains concluding remarks.

The basic result is that the classical principles of demand and supply can be formulated in two natural ways to understand the price patterns that are observed. A model of a *temporal equilibrium* (TE) is developed and is shown to closely mirror the complex price patterns that result from the random arrivals and departures inherent in the

² Partial and special case models have been developed to study specific behaviors found in traditional experimental environments. Such models are stylized and would require theoretical generalization to be usefully applied to the continuous market environment we explore here. Suggestive theoretical developments do exist. See for example, Parlour (1995) and Goettler, Parlour and Rajan (2005) who focus on the strategic structure of bids and asks in a theoretical finance setting. Their theoretical environment can be interpreted as a special case of the experimental environment developed here.

environment. A second model called the *flow competitive equilibrium* (FCE,) representing the implications of the flow characteristics of market incentives, is also developed and captures the underlying probabilistic structure of incentives and arrivals. The FCE can also be used to model the behavior of prices when averaged over long ranges. We discover that these two equilibria are interactive, exerting independent influences on prices. The markets demonstrate high levels of efficiency, indicating that, despite a lack of direct coordination of market timing, the potential gains from trade do tend to evolve.

2. THE RANDOM ARRIVAL AND DEPARTURE ENVIRONMENT

The new environment has many similarities with classical environments. Both environments have agents with preferences, the difference being that in the random arrival environment the preferences are characterized as probabilistic. We use the term “latent” preferences to refer to the new concept. The classical environment typically has fixed agents or agents that arrive as a group at the same time while the random arrival environment has incentivized agents arriving according to some stochastic process in time

Preference Inducement Methodology

Classically experimental market environments, as developed by Chamberlin (1948) and Smith (1962), consists of a set of redemption values, costs, and a period structure. Before the start of a period, buyers receive redemption values from the experimenter and sellers receive costs. Buyers make money in an experiment by buying units in a public market, in which all subjects can participate, and reselling them to the

experimenter at the redemption values the experimenter privately quotes each buyer. Similarly, sellers buy units from the experimenter, at costs the experimenter quotes, and resell them to other subjects for a profit. Under the assumption of the competitive model, redemption values and costs can be modeled as limit prices and used as parameters in a market model of competitive supply and demand equilibrium. When a period opens, subjects choose what incentives they will act on and form trades in the public market. Each period typically lasts for a fixed length of time. After each period, subjects receive additional redemption values and costs, while old redemption values and costs do not carry forward to new periods. Additionally, units that exist in one period typically are not carried over to the next period; inventories and cash typically refresh each period.

In the classical environment, each period is like a day in which commodities are traded and completely depreciate over night. The day starts with a stock of costs and redemption values. During the day the gains from exchange explicit in the stock are exhausted. All actions are coordinated by the beginning and ending of the period.

By contrast, the random arrival environment has no period structure. The market opens for a fixed length of time, typically approximately two or three hours. Incentives arrive in the form of private orders to buy from the experimenter (i.e. costs for potential sellers) or private orders to sell to the experimenter (i.e. redemption values for potential buyers) in a market accessible only by the agent for whom the orders are intended (i.e. the agent's private market)³. Buyers have an opportunity to buy in the public market from other agents and resell for a profit in their private market by accepting an order to sell to the experimenter found there. Similarly, sellers accept private orders to buy from

³ This method of implementing the random arrival of incentives is made possible by the Caltech Marketscape technology that will be explained in greater detail in later sections.

the experimenter found in their private markets and resell units to other agents in the public market.

Private orders to buy and sell appear in agents' private markets at random arrival times and each order will expire after a short period if not acted on. This expiration feature is important because it forces the individual to decide whether or not to act on an order during a specific interval of time. The incentives can appear at any time for any subject and last as long as the experimenter chooses. Thus, at any instant, a subject can have several orders for different amounts that appeared in the subject's private order book at different times and have different expiration times.

Incentive Parameter Structure (Latent Incentives and Realized Incentives)

The basic parameters will be called “*latent buyer incentives*” and “*latent seller incentives*”. The latent buyer incentives consist of a probability density function $g_b(x)$, where x is a price. Latent seller incentives consist of a probability density function $g_s(y)$, where y is a price. For individual agents, draws are made from the distribution of buyer values and the distribution of seller costs according to two independent Poisson processes with intensities λ_s and λ_b respectively.

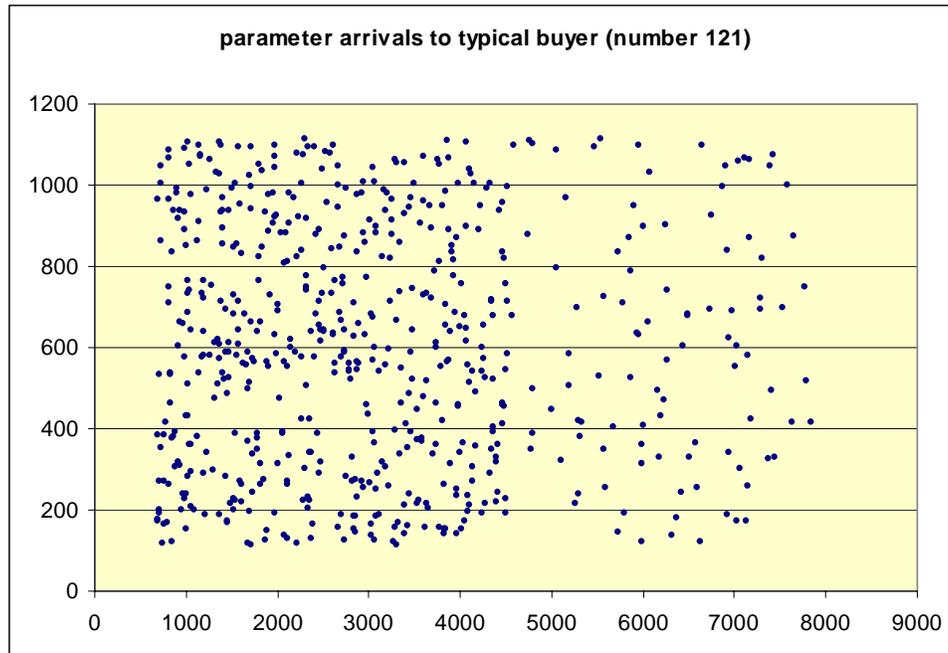
Realized incentives, as opposed to latent incentives, are the draws that are actually sent to buyers' and sellers' private order books and serve as “redemption values” and “costs”. In designing experiments, λ_s is the arrival rate of private orders for each of the n_s sellers, and λ_b is the arrival rate of private orders for each of the n_b buyers. Here we assume that all agents of a type receive incentives at the same rate. An order sent to a private order book has a life δ_b and δ_s for buyers and sellers respectively. In these

experiments, δ_b and δ_s are fixed lengths of time (6 minutes), but this need not be true in general. The environment could easily be modified to include random expiration according to some waiting time distribution.

One can think of nature randomly choosing buyers at a rate $n_b \lambda_b$ from a distribution g_b of latent buyer types with each type being a person's willingness to pay. Similarly for sellers. Thus, we will sometimes say loosely that the buyers and sellers are randomly arriving at the market with randomly distributed incentives and a fixed life.

Figure 1 provides an impression of the environment from the point of view of a subject. Shown there are realized incentives (the private orders received) by a subject over the course of an experiment. The horizontal axis is the time of arrival and the vertical axis is the price of the private offer (the analog of a "redemption value"). A parameter shift to a lower arrival rate took place about the middle of the experiment. As can be seen from the pattern, the subject faces a wide range of randomly arriving incentives. When all signals are viewed at once, as is the case in the figure, the difference in the pattern of incentives is apparent. Things are more subtle from the subject's point of view. Only the arrivals themselves are observed by the subject without aggregation or frequency measurements. The subject is only exposed to a change in the arrival rate and this change is not signaled by other features of the environment.

Figure 1: Example Arrival of Private Orders (Incentives) Before and After a Parameter Shift That Reduces the Flow of Orders to the Subject.



While the environment introduced here is new, the experimental literature contains suggestive departures from the classical environment. The literature is much too large for a complete review here, so, we do not attempt to review all of the modifications of the classical environment that exist in the experimental literature. Instead we reference seminal, major departures in the direction of the environment developed here.

In Jamison and Plott (1997) and Kagel (2004) the incentives differed each period in a random fashion, In Brewer *et al* (2002), incentives were instantaneously refreshed after a trade took place, demonstrating that the price adjustment process was not due to the Marshallian path.⁴ Many experiments involve incentives with multi-period

⁴ Interestingly, because the units that could trade would be refreshed, the “arrival” rate of such units began to increase relative to units that could not trade. In order to accommodate this feature, Brewer, *et al* (2002)

longevities following the original study by Miller, Plott and Smith (1977): notable examples being experiments with financial assets (Foresythe, Palfrey, and Plott, 1982; Smith, Suchanek and Williams, 1988) and many other experiments involving goods with “asset-like” properties (Lian and Plott, 1998).⁵

A flow environment with simulated buyers was created by Millner, Pratt, and Reilly (1990) for the study of contestable markets characterized by duopolists with falling average cost, but they studied only a solution from contestable market theory as opposed to a general concept of competitive market equilibrium. Aliprantis et. al., introduced the idea of “overlapping orders” similar to the idea of “overlapping generations” which have features similar to the random arrival markets we introduce here.⁶ In the overlapping orders environment, each agent-type had a fixed period structure, say every 20 minutes, the beginning of which orders arrived that could be executed during the personal period and expired at the end of the personal period. Same agent types operated on the same schedule with essentially identical preferences while different agent types operated on different (overlapping) period schedules. For example, in a two generation world, the periods for generation 2 started 10 minutes after the period for generation 1 started. The market never closed so at each instant there was a “young generation” that just received incentives and an old generation, with incentives that were getting ready to expire. Thus, the classical period structure was removed. One can think of the random arrival

invented a “velocity adjusted” concept of demand and supply that can be viewed as a special case of the theoretical concepts developed here.

⁵ Cliff and Priest, (1998) allowed the accumulation of inventories and orders that were distributed to subjects at scheduled periods.

⁶ In a much different environment, overlapping generations have been studied by Marimon and Sunder (1993).

environment as an “overlapping order” environment only with random schedules that differ across individuals and many generations.

3. MARKET INSTITUTIONS

The market organization implemented here is the multiple unit double auction with an order book invented by Plott and Gray (1990). At any instant, a buyer or a seller can submit an order consisting of a quantity, a per-unit price and an expiration time and send it to the market. Buy orders obligate the bidder to buy up to the stated quantity at the per unit price. Sell orders obligates the asker to sell up to the stated quantity at the per unit price. Orders are sent to a public order book that can be viewed by all agents and are listed in order of price from best to worst from the point of view of counterparties.

If trade is possible when an order arrives at the market, the trade is immediately executed at the existing price in the order book. That is, if a buy order arrives at a price that is higher than the lowest sell order price the trade is executed at the sell order price. If the quantity of either side is not exhausted the remaining amount is entered into the book.

The market exchange system was Caltech’s Marketscape program. This market system operates over the web; agents can be located at different institutions or at home. The exchange system has a public market in which exchanges can take place. Each agent also has a private market in which orders are placed by the experimenter. These private markets provide the technology through which the random arrival environment is implemented.

4. MODELS AND THEORY

Classical Models

The models we propose for application in the random arrival environment are closely related to the classical principles. These properties, as related to the basic principles, are as follow:

1. The law of one price in a market. This idea is fundamental and serves as a fictional concept of “market price.” Classical experiments demonstrate it tends to be more than a fiction since the variance of trade prices falls over periods in experiments with a stationary environment across periods. While a single, constant price is seldom observed, the prices are “close together”. This property is thought to result from arbitrage within a period. Individuals are able to buy and sell within a period and force prices together. Of course, notable exceptions are the bubbles in asset markets.
2. The competitive equilibrium is a (single) price that equates the competitive quantity demanded with the competitive quantity supplied in the model (intersection in the case of correspondences). With replications of periods observed prices move close to the equilibrium. Exactly why this occurs is unknown and has been the subject of theory Easley and Ledyard (1993), Wilson (1987), Friedman (1991), Satterwaite and Williams (1993), and Asparouhova, Bossaerts, and Plott (2003).
3. The competitive market demand and supply are defined by the aggregation of individual limit prices (where limit prices are understood in the sense of inverses of individual demand and supply functions). The limit prices are

assumed to be the redemption values and costs (private orders in our terminology).

Flow Competitive Equilibrium Vs Temporal Equilibrium

Two natural models of equilibrium suggest themselves. The temporal model captures the state of the market at an instant of time. The flow competitive model captures the state of a market in a probabilistic sense. Both models evolve from the “as if” assumption of a "market price" taken as a constant by participants. Market supply and demand are defined as having resulted from an individual optimization subject to opportunities and beliefs. The temporal equilibrium is calculated using incentives that exist at a given moment, while the flow model is probabilistic.

Temporal Equilibrium

At any given time, *temporal competitive supply* (TS) and *temporal competitive demand* (TD) curves are based on orders that exist in private order books (private incentives) at time t . These are the orders received by subjects that have not been acted upon or expired. For subjects i and j let $R^i(t, x_i)$ be the revenue that is produced by exercising the best x_i orders that buyer i finds in the private order book at time t and let $C^j(t, y_j)$ be the cost of buying the best y_j order found in seller j 's order book at time t . Let P be the market price. The temporal competitive model holds that x_i is chosen to $\text{Max} [R^i(t, x_i) - Px_i]$ and y_j is chosen to $\text{Max} [Py_j - C^j(t, y_j)]$. From the optimization model the TD and TS are always well defined for the individuals and the TD and TS are well

defined at the market level as the sum of the functions for the individuals at a given market price.

From the construction above, we know that the temporal demand curve at time t is a downward sloping step function, $TD(P,t)$, equal to the number of buyers (sell orders in private markets) in the market at time t —those that have arrived before t and have not yet either traded or were cancelled—with reservation prices above P . Similarly $TS(P,t)$ is an upward sloping step function equal to the number of sellers (buy orders received in private markets) with reservation prices below P at time t . We can define a temporal equilibrium price, as a P such that $TD(P,t) = TS(P,t)$.

Flow Competitive Equilibrium

Flow competitive demand (FCD) and *flow competitive supply* (FCS) curves, on the other hand, specify the arrival rates of buyers (sellers) with reserves above (below) a given price. Flow competitive supply and flow competitive demand reflect two components: 1) the distribution of latent reservation prices for buyers and sellers, and 2) the relative arrival rates of buyers and sellers. For a given price P , the levels of the flow competitive supply and demand curves are given by:

$$\begin{aligned} FCS(P) &= n_s \lambda_s \int_{-\infty}^P g_s(y) dy = n_s \lambda_s G_s(P) \\ FCD(P) &= n_b \lambda_b \int_P^{\infty} g_b(x) dx = n_b \lambda_b (1 - G_b(P)) \end{aligned} \quad (1)$$

Where λ_s is the arrival rate of individual sellers, λ_b the arrival rate of individual buyers, n_s and n_b are the number of seller-participants and buyer-participants, and g_s and g_b are the latent preferences, the distributions of reserve prices for sellers and buyers respectively.

A *flow competitive equilibrium* (FCE) is defined by 1) a price P at which the arrival rate of buyers with reservation prices at or above P is equal to the arrival rate of sellers with reserve prices at or below P and 2) a rate of trade associated with P . That is, the FCE is a price, P_e , and flow competitive equilibrium transaction rate λ_{FCE} defined by:

$$\begin{aligned} FCS(P_e) &= n_s \lambda_s G_s(P) = n_b \lambda_b (1 - G_b(P)) = FCD(P_e) \\ \lambda_{FCE} &= n_b \lambda_b \int_{P_e}^{\infty} g_b(x) dx \end{aligned} \quad (2)$$

The FCE price is the price such that the flow of supply equals the flow of demand. The equilibrium flow is simply the FCD evaluated at the FCE price.⁷

Figures 2a and 2b illustrate graphs of FCS and FCD produced from uniform distributions of reserve prices on 0 to 1000. Figure 2b, shows how the curves in 2a change with changes when the rate of arrival for buyers is cut in half, while 2c shows how FCS and FCD change when the distribution of buyers' valuations is shifted upward. Figure 2d illustrates how the FCS and FCD curves generalize to different distributions of

⁷ The FCE can be viewed as a connection with theoretical ideas under development in finance. Close relationships exist between the environment introduced here and the theoretical financial market explored by Goettler, Parlour and Rajan (2005). In a sense, their environment can be viewed as a special case of ours. The prominent features of their environment are: (i) private values that “reflect the idiosyncratic motives for trade (wealth shocks, tax exposure, hedging, or portfolio rebalancing needs)”; (ii) the independent arrivals of traders drawn from known distributions; (iii) a publicly known “consensus value” of an asset, perhaps dictated by the present value of a dividend stream; and (iv) upon arrival in the market the trader makes a decision about the type of order to place in an open order book and implicitly, the timing.

The essence of (i), (ii), and (iv) are in both ours and the GPR environment. A concept of a “consensus value” as found in (iii), can be found in both, but in the environment introduced here, it emerges as a candidate equilibrium concept, the FCE, as opposed to an imposed parameter as done in GPR. While the FCE carries much of the intuition carried by the “consensus value” of GPR, it is not public information and there are both conceptual and technical differences. For example, when buyers and sellers have a common distribution of latent preferences and the arrival rates are the same, the FCE is the median of that distribution while the consensus value of GPR would be the mean. In addition, the FCE generalizes to the cases where the latent preferences of potential buyers and sellers do not arise from a common distribution and, since the FCE is closely associated with the classical competitive model, information or common knowledge about underlying parameters play no particular role.

incentives by using a truncated normal distributions with a mean on 500f and a variance of 200f to generate the curves.

Figure 2a: Flow Competitive Supply and Demand Arrival Curves with 1000 Buyer and Seller Arrivals Per Hour.

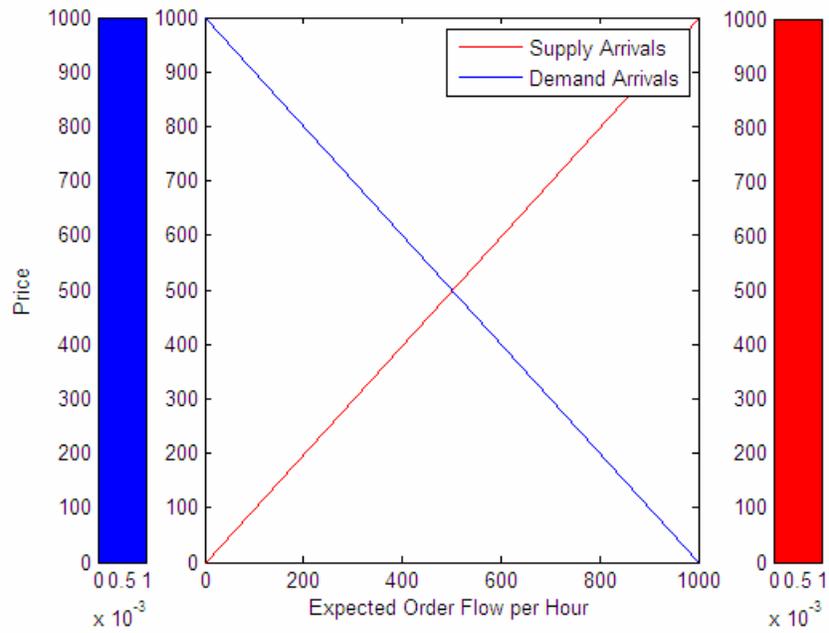


Figure 2b: Flow Competitive Supply and Demand Arrival Curves with 500 Buyer and 1000 Seller Arrivals Per Hour.

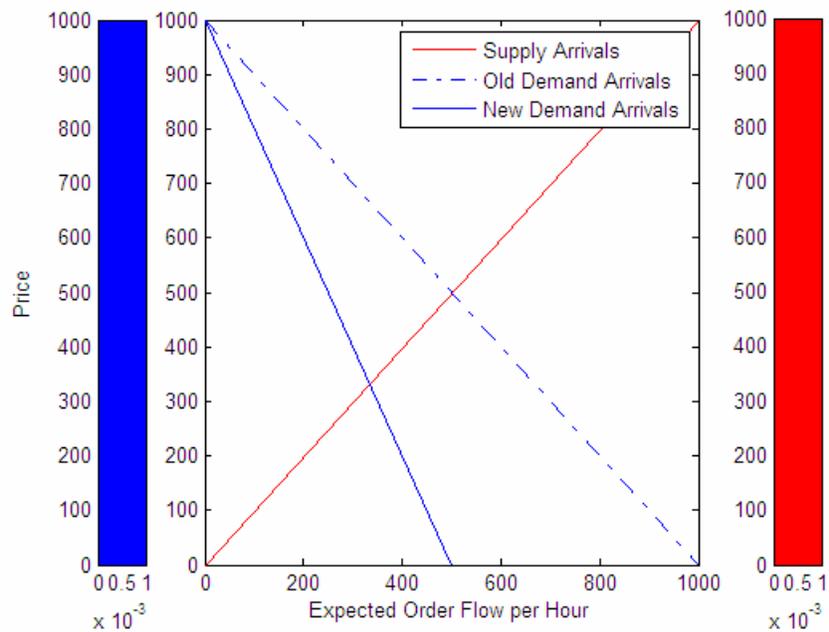


Figure 2c: Flow Competitive Supply and Demand Arrival Curves with 1000 Buyer and Seller Arrivals per Hour and Shifted Latent Demand.

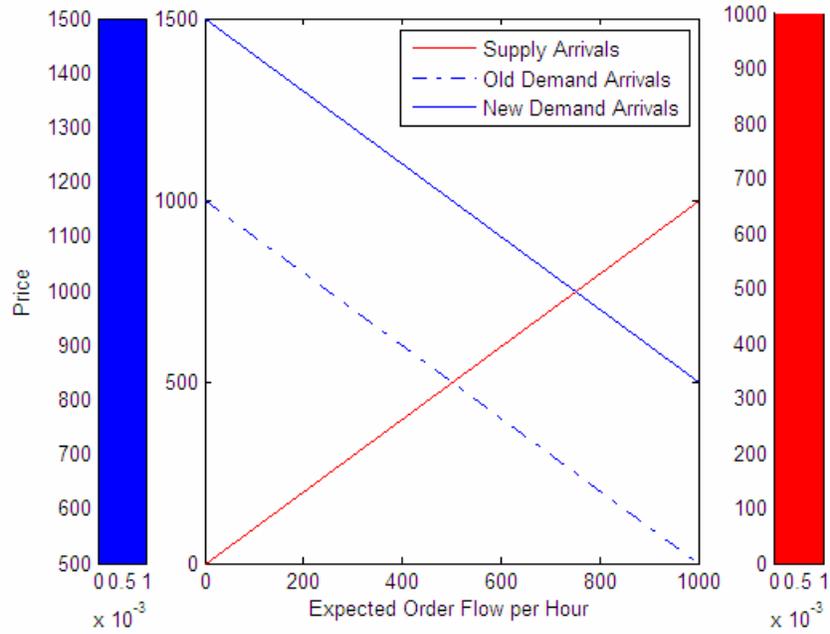
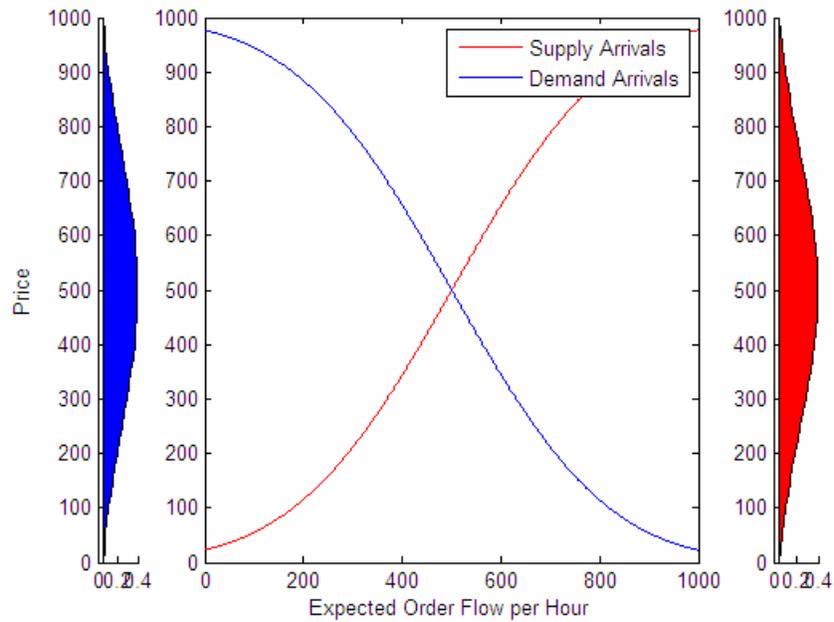


Figure 2d: Flow Competitive Supply and Demand Arrival Curves with 1000 Buyer and Seller Arrivals per Hour and Normally Distributed Latent Incentives.



One way to think about the relationship between temporal and flow competitive equilibrium concepts is the following. Temporal supply and demand curves are essentially *empirical* cdf's which have been turned on their sides and stretched by some factor. Flow competitive supply and demand curves, on the other hand, are essentially *parametric* cdf's which have been similarly turned on their sides and stretched by some factor. A property of empirical cdf's based on n observations is that as n goes to infinity, the empirical distribution converges point-wise to the parametric distribution. Hence, temporal supply and demand curves converge to competitive supply and demand curves as the number of orders in the market becomes large. The number of orders in a market can become large for many different reasons, such as aggregation of orders over time, an increase in arrival rates, or an increase in the longevity of orders, etc.

It is easy to show, using the logic above, that as the number of incentives on both sides of the market goes to infinity, the expectation of TE prices converges to the FCE price. Hence, the TE price is an asymptotically consistent estimator of the FCE price. Given consistency (a large sample property), a natural question is whether the TE is also an unbiased estimator of the FCE for finite numbers of buyers and sellers in the market (a small sample property).

We provide no proof of the unbiasedness of TE, but conjecture that this property is true for some set of “nice” latent preference distributions. We test this property for each of the parameter sets used in our experiments using simulations and find that it appears to be true for the parameters considered. Two similar properties of TE prices are listed below together with sketch proofs of each statement.

1. Let $ED_w(P)$ represent the excess demand that results from accumulating w minutes of order flow. If P^* is a unique FCE price, then P^* is the unique price P at which $E[ED_w(P)]=0$. More generally, any price at which the arrival rates are equal implies that expected excess demand is zero.⁸
2. Suppose that over a period of length w , n incentives sequentially arrive to the market. Suppose also that by the time the $n-1$ th incentive arrives there exists a well defined TE price defined as the midpoint of an incentive crossing. Let $TEQ(n)$ and $TEQ(n-1)$ be the TE that obtains at the time of the n th and $n-1$ th arrival.

$$E[TEQ(n)|TEQ(n-1)<P^*] \geq TEQ(n-1), \text{ and}$$

$$E[TEQ(n)|TEQ(n-1)>P^*] \leq TEQ(n-1)^9$$

The following two figures illustrate the simulated relationship between TE and FCE. Figure 3a shows that the distribution of TE prices is on average equal to the FCE price. Shown there is the distribution of TE prices for a market with an arrival rate of 8 buyers and sellers per 30 seconds and a distribution of reservation prices uniform on 1 to 100 francs as a function of how long orders are aggregated over time. As w , the length of the observation window, becomes large, the distribution of TE prices becomes tight around the FCE price of 50.5 francs.

⁸ Sketch of proof: $E[ED_w(P^*)]=E[D_w(P^*)]-E[S_w(P^*)]=\lambda_{FCE} * w - \lambda_{FCE} * w = 0$
 $E[ED_w(P')] = E[D_w(P')] - E[S_w(P')]$, $E[D_w(P')] \neq E[S_w(P')]$ by uniqueness

⁹ Sketch of proof: $TEQ(n)$ will tend to move up (down) when new buyers (sellers) with limit prices above (below) the current TE arrive to the market. If the current TE price is below (above) the FCE price, then the arrival rates of buyers with reservation prices above the current TE price will be faster (slower) than the arrival rate of sellers with costs below (above) it. Hence, there is more probability that the next arrival will be an incentive which shifts the new TE price above (below) its current value.

The second figure, Figure 3b, illustrates that the property extends itself to the case in which the FCS and FCD are asymmetric. That is, the asymmetry between the price generating curves does not result in an asymmetry in the distribution of prices as one might expect from studies that show the distribution of consumer and producer surplus generates an asymmetry in the distribution of transaction prices (Smith and Williams, 1983). The figure shows the same information, but for a market that is asymmetric. Here buyers arrive at a rate of 4 per 30 seconds, while sellers arrive at a rate of 16 per 30 seconds. This pushes down the FCE price to around 20 francs. Again, the distribution of TE prices is centered on the FCE price and becomes more tightly distributed as we aggregate more orders over time.

Figure 3a: Simulated Distribution of TE Prices.

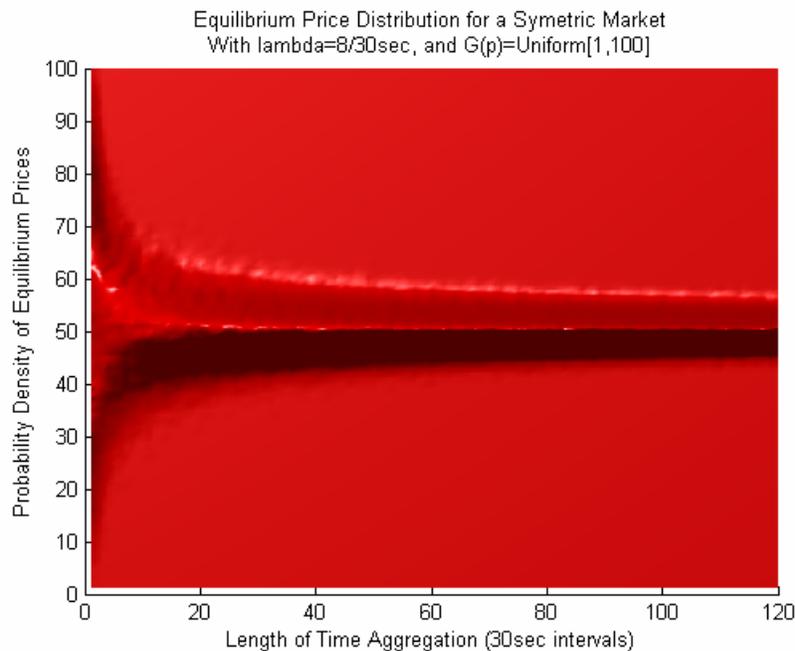
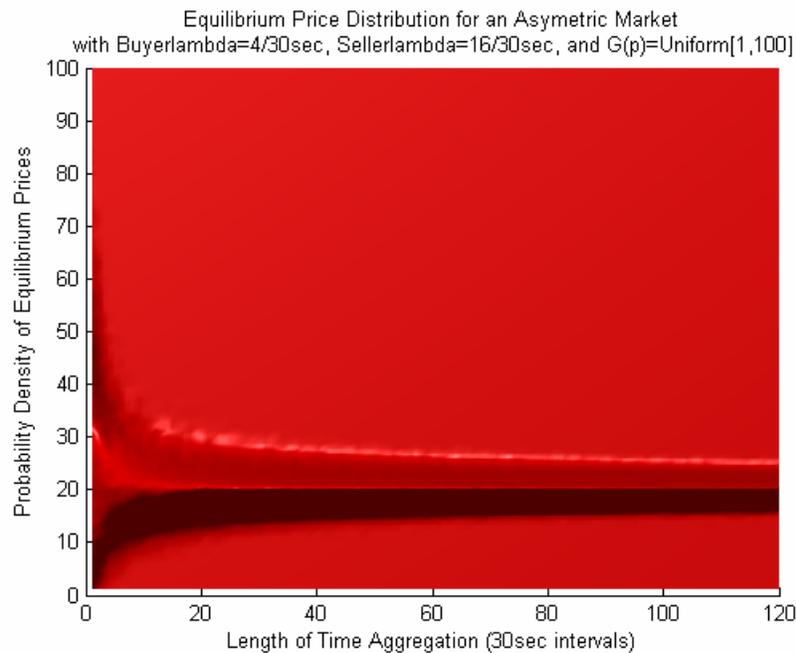


Figure 3b: Simulated Distribution of TE Prices.



5. EXPERIMENTAL PROCEDURES AND DESIGN

Experimental Procedures

Subjects were students recruited from Claremont McKenna College, Occidental College, and Caltech by a general request for people to put themselves in a database if they were interested in participating in experiments. The day before an experiment, invitations were sent via e-mail recruiting subjects from that database. Typically, these experiments recruited subjects from more than one school.

Subjects who reserved a spot in an experiment were sent the web location of a training program that allowed them to participate as buyers and sellers in market software typical of the market mechanism to be used in the experiment. Several of the students, especially those from Caltech, had prior experience with economics experiments in general. A few subjects had prior experience with market experiments in particular.

Subjects were asked not to reserve a spot in experiments unless they were able to show up and participate in the whole experiment, but nearly every experiment had either subjects that were “no-shows,” or subjects that dropped out before the end of the experiment. Experiments were conducted either in the evening, (around 7:00PM) or on weekends.

Subjects were given the web address of the experiment and told that they could go to the web address to get an identification number and password. Instructions were also posted at the experiment location. These instructions were very similar to the instructions that existed in the practice program except for the amounts of incentives and the timing when they might arrive.

Each experiment was preceded by a ten minute practice period for which subjects did not receive payment. The practice parameters were unrelated to those used in the experiment. Subjects’ trading activity was monitored remotely to determine whether subjects were confused about whether they were a buyer or seller, or were confused regarding how to use their private markets. Subjects were additionally provided a phone number that they could call with any questions they had about the experiment.

The experiments started on time. At the end of the experiment subjects were told to check their mailing addresses in the database and to check our calculation of how much they earned. They were sent a check for their earnings. Subjects earned anywhere from \$10 to \$78 for a two hour experiment depending on performance, with most subjects earning close to an average of \$40.

Experimental Design

A total of five experiments were conducted.¹⁰ Each experiment featured one shift in either the distribution of buyers' redemption values/sellers' costs or a shift in the rates of arrivals. The times of these shifts occurred near the middle of each experiment and are recorded in Table 1. Also recorded in Table 1 is the length of the experiment, the number of buyers and sellers, the total number of incentives sent to buyers and sellers before and after the shift, as well as the distributions of incentives and the FCE before and after the shift.

In designing the experiments, order-flow parameter files were constructed on a per person basis according to a Poisson process with redemption values/costs drawn independently from distributions known to the experimenters but not to subjects. Because of this, the experimenters did not know the actual numbers of incentives that would arrive on the buy and sell sides of the market in advance. For each buyer and seller, the experimenters recorded the time of their first and last action in their private market. The number of incentives sent to the market listed in Table 1 includes only those incentives that were in the market, or arrived to the market, during the interval that the trader for whom they were intended was active.

Since the experiments were conducted with remote subjects, tight control over participation was impossible. Parametric adjustments to models were required when subjects quit the experiment after having started. In such cases the models were recalibrated for a different number of subjects beginning from the time that the subject stopped participating. For most experiments, the adjustment made for when traders were present in the market was not important. Only in experiment market070414, were there

¹⁰ An additional four experiments were run as pilots but were not included in this study due to the choice of parameters, computer problems during the experiment, or small sample sizes.

drop-outs and late entrants which affected the calculation of FCE. These all occurred before the parameter shift and will be illustrated in the figure that plots the FCE price path for this experiment.

Table 1: Summary of Experiment Parameters

ExpDate	Experience	Average Earnings	School(s)	Buyers	Sellers	Number of Buyer Incentives	Number of Seller Incentives	Buyer Arrival Rate	Seller Arrival Rate	Buyer Distribution	Seller Distribution	Exp Length (min)
70208	Inexperienced	approx \$40	CMC	8	6	(1520, 2225)	(1124, 1607)	1/15sec, 1/15sec	1/15sec, 1/15sec	U(273,672), U(52,451)	U(273,672), U(52,451)	129.85
70414	Mixed, Moderately Experienced	approx \$40	CMC,Oxy	9*	9*	(4315, 2344)	(4701, 3292)	4/15sec, 2/15sec	4/15sec, 4/15sec	U(631,1632), U(631,1632)	U(631,1632), U(631,1632)	127.23
70420	Mixed, Moderately Experienced	approx \$40	Caltech, CMC, Oxy	8	9	(3741, 748)	(863, 2480)	17/8sec, 1/2sec	1/2sec, 17/8sec	U(114,1115), U(114,1115)	U(114,1115), U(114,1115)	127.23
70425	Mixed, Mostly Experienced	approx \$40	CMC, Oxy	7	8	(830, 3302)	(3633, 941)	1/2sec, 17/8sec	17/8sec, 1/2sec	U(228,1229), U(228,1229)	U(228,1229), U(228,1229)	135.05
70606	Mixed, Moderately Experienced	approx \$40	Caltech, Oxy	7	5	(1848, 1122)	(1106, 773)	1/15sec, 1/15sec	1/15sec, 1/15sec	U(52,451), U(273,672)	U(52,451), U(273,672)	119.80

*includes dropouts

6. RESULTS

The results are divided into three parts. The first part is a simple overview of the experiments containing figures of price patterns and parameters. The second part focuses on efficiency of the market. The third part focuses on price levels.

Overview

Perhaps the best way to form an impression of the data and the results is from the graphs of parameters, models and trade prices. Figures 4a thru Figure 8b provide an overview of the nature of the price data and results for all experiments. For each experiment the “a” parts of the figure plots the incentives forming supply and demand before and after the shift. The “b” parts of the figures show the time series of (i) the FCE, (ii) the TE price and (iii) trade prices.

As will be stated in the efficiency section that follows, these markets exhibit very high levels of efficiency. The reason for the high efficiency levels can be deduced from two facts that can be seen in the figures and will be stated precisely in the section of pricing results. First, trade prices are closely related to both TE and FCE prices; with TE prices changing rapidly over the course of an experiment. Second, averaged over the entire experiment, the TE prices and the trade prices are near the FCE. As is clear from the figures, trade prices exhibit considerable variability. Overall, the new theories of supply and demand developed above do a good job of telling us what to expect in such complex markets.

Figure 4a: Flow Competitive Supply and Demand Parameters for Market 070208

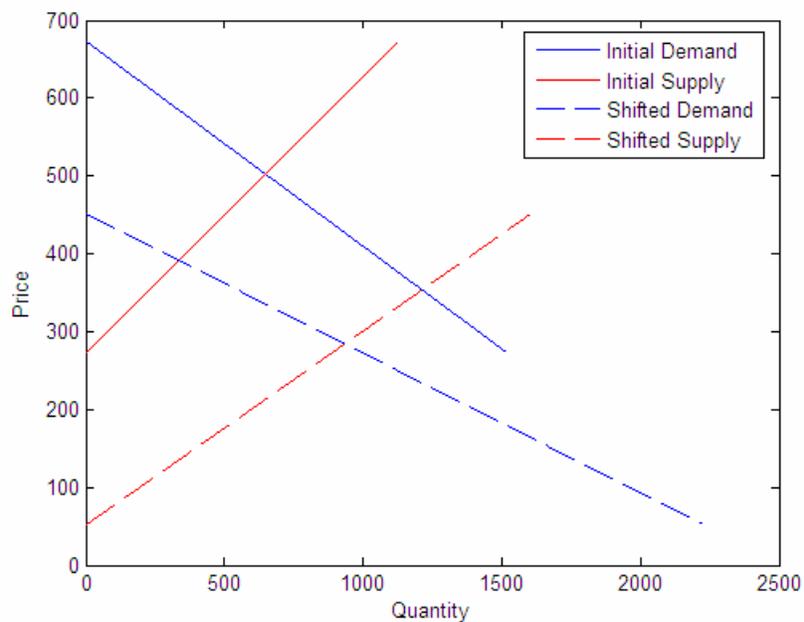


Figure 4b: Experimental Results of Market 070208

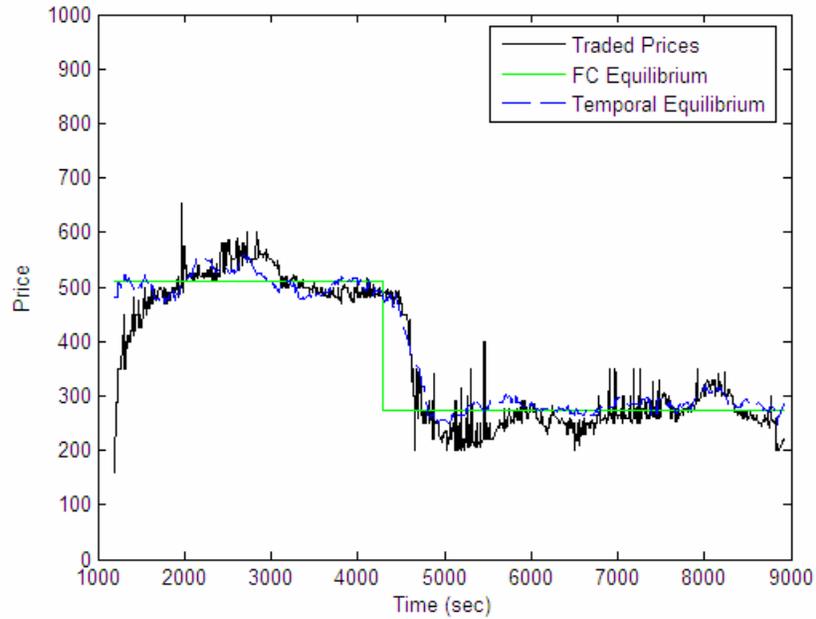


Figure 5a: Flow Competitive Supply and Demand Parameters for Market 070414

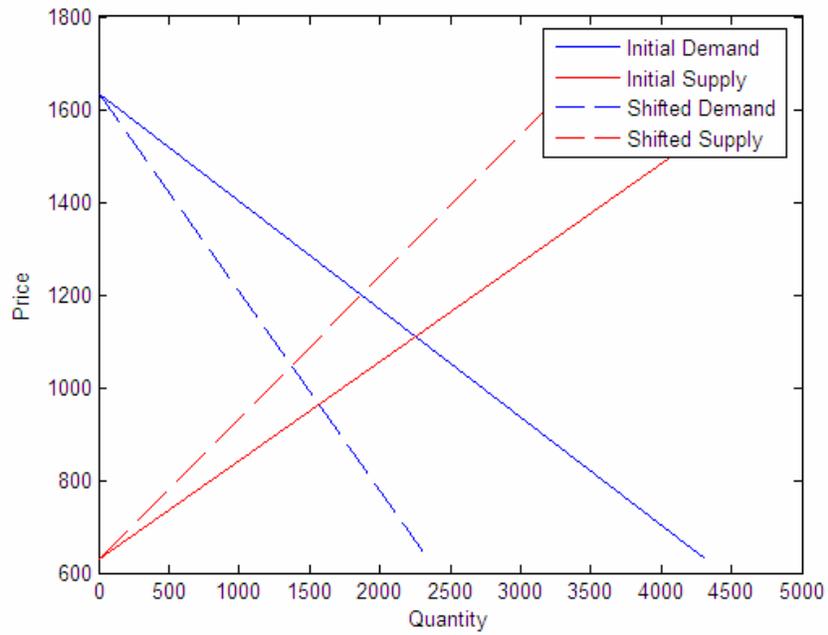


Figure 5b: Experimental Results of Market 070414

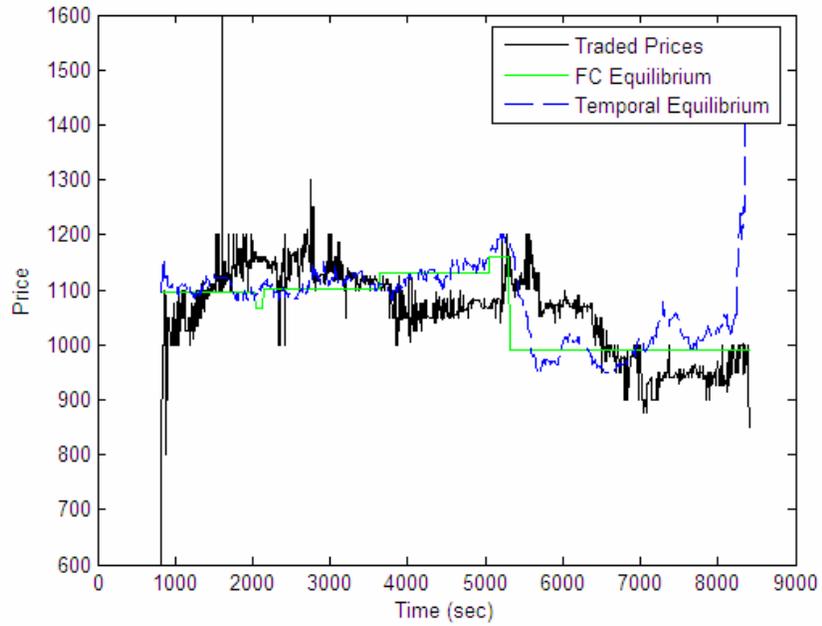


Figure 6a: Flow Competitive Supply and Demand Parameters for Market 070420

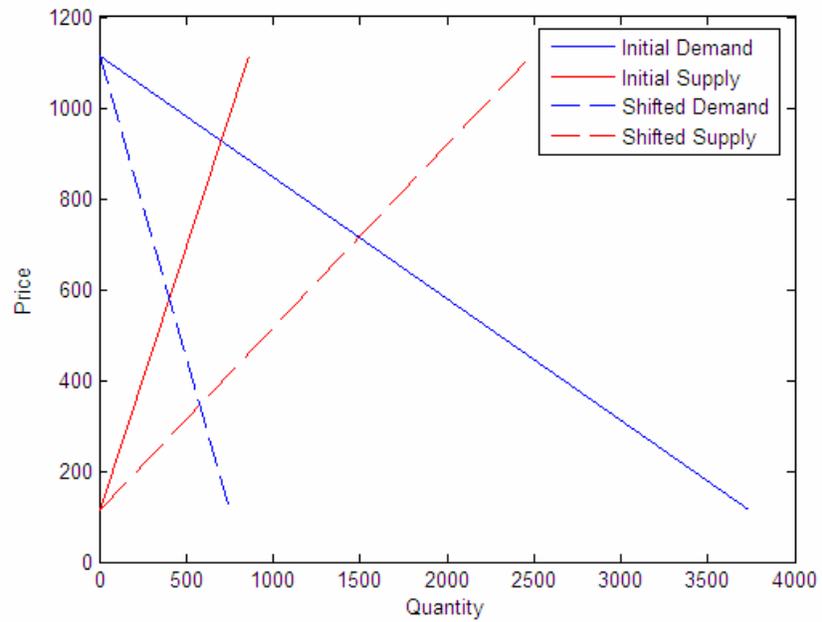


Figure 6b: Experimental Results of Market 070420

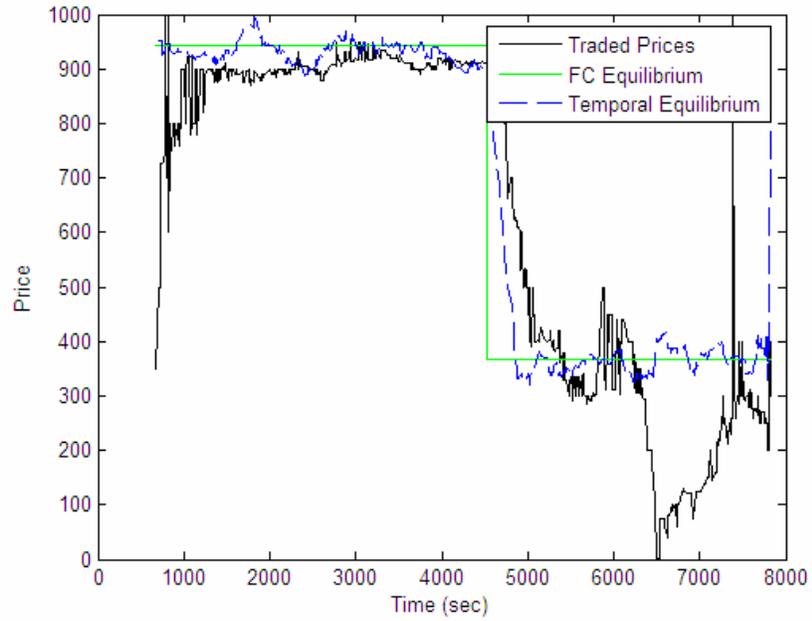


Figure 7a: Flow Competitive Supply and Demand Parameters for Market 070425

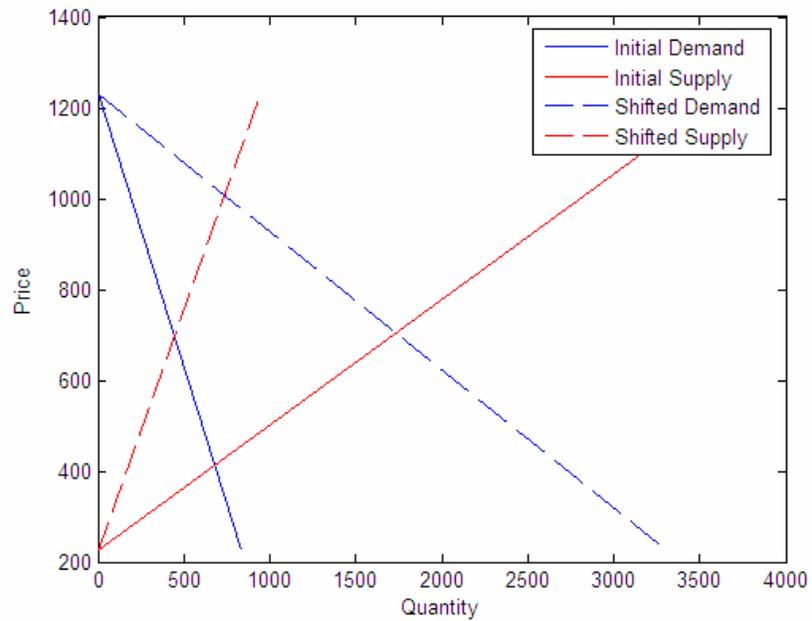


Figure 7b: Experimental Results of Market 070425

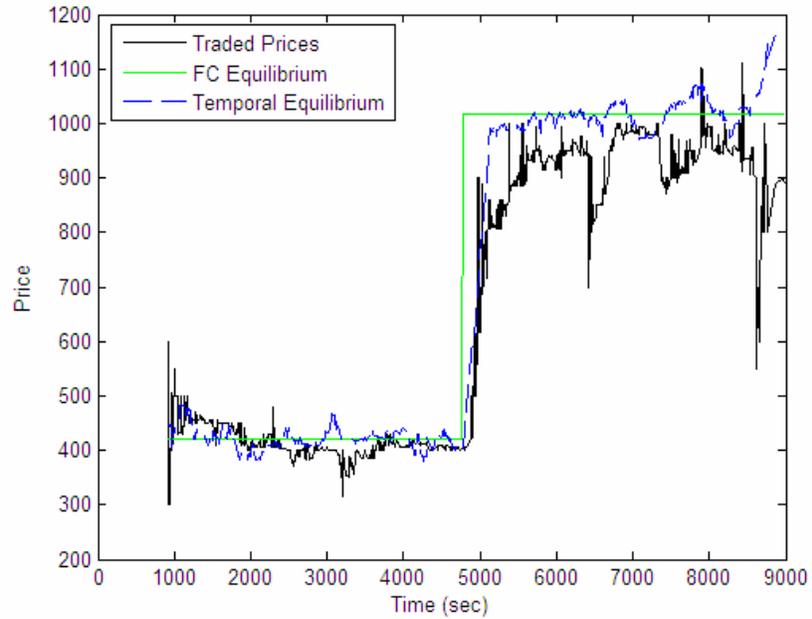


Figure 8a: Flow Competitive Supply and Demand Parameters for Market 070606

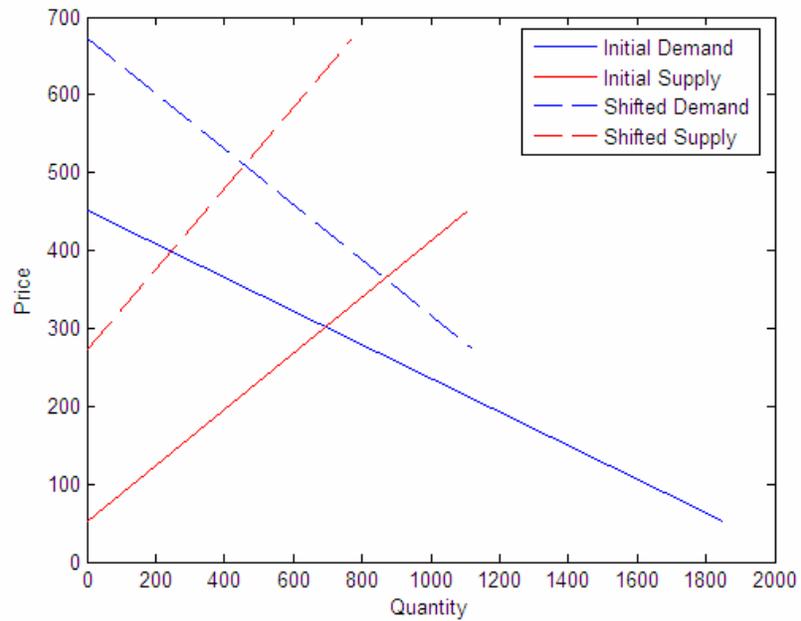
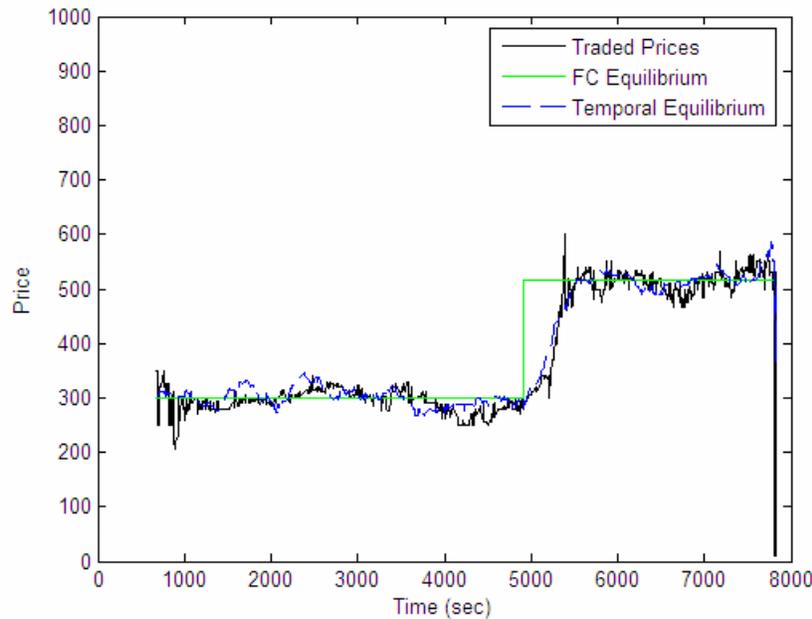


Figure 8b: Experimental Results of Market 070606



Efficiency

In an environment with incentives arriving at different times there can be multiple different definitions of efficiency. Of course, each efficiency concept is closely related to the concept of experimental market efficiency first developed by Plott and Smith (1978). Table 2 reports the efficiency of each experiment relative to three different efficiency measures. The first measure compares the total gains from trade to the maximum possible gains from trade. In essence, this is the surplus that would be obtained if the all incentives, before and after the shift were aggregated as a stock, a single FCE price solved for, and all trades occurred at that price. We will refer to this fraction of the maximum surplus attainable as the *clairvoyant efficiency level*, because in order for a trading mechanism to attain the maximum possible surplus, it would require a

foreknowledge of future incentives flow and parameter shifts. This measure of efficiency will always be less than one.

The second efficiency concept compares actual trading surplus to the level that would be obtained if all trades involving incentives that arrived prior to the shift occurred at the initial FCE price, and all of the trades involving incentives which arrived after the shift occurred at the second FCE price. We refer to this as *flow competitive rational efficiency*. This efficiency measure is not necessarily between zero and one.

The third level of efficiency reported, *local incentive efficiency*, compares actual surplus with the amount that would be obtained if traders submitted bids and asks equal to their reservation prices immediately when they received an incentive. Under this trading strategy, there are no gains from trade due to price smoothing, or speculation, which allows gains from trade to be realized between two traders who are not in the market at the same time. This compares actual surplus to a particularly low level, and so it is not surprising that this efficiency measure is typically above 100%. The amount by which this measure exceeds 100% can be interpreted as the amount of surplus traders gained by smoothing prices over time.

Result 1: Trading in experimental flow markets generates high levels of efficiency relative to the maximum amount of surplus available. Realized surplus extraction is typically higher than the local incentive efficiency that could be obtained without smoothing/speculation.

Support: Table 2 shows that the levels of efficiency relative to each measure are remarkably high, even when compared to the maximum attainable surplus. The flow competitive rational efficiency is near 100% and the local efficiency exceeds 100%. These measures suggest the existence of efficiency gains due to speculation over time.

Table 2:

ExpDate	Clairvoyant Efficiency	Flow Competitive Rational Efficiency	Local Incentive Efficiency	Trades Actual/Predicted
70208	75.740%	91.710%	135.93%	1878/1582**
70414	87.230%	87.370%	125.41%	4908/3596**
70420	64.250%	96.000%	100.10%	1713/1281**
70425	60.950%	94.120%	99.12%	1824/1407**
70606	90.600%	102.050%	135.61%	1458/1114**
**Ratio reflects speculative trades not included in order flow parameters				

For the local efficiency to exceed 100% one suspects speculation in the absence of some special coordinating device. That observation, in turn, suggests that volume be examined. The next result demonstrates that no special coordination is obvious and that the volume is higher than predicted by the FCE.

Given that incentives arrive to the market according to a Poisson process, a reasonable and testable hypothesis is that the process of trading will also occur according to a Poisson process, possibly dependent on recent order flow. The following result provides an overview of the trading volume pattern and demonstrates that it is not closely constrained by the underlying flow of incentive arrivals.

Result 2: Waiting times between trades are uncorrelated, and have a mean rate of transaction larger than the rate of transaction predicted by the FCE.

Support: A part of this result can be seen in Table 2, which lists the total number of transactions for each experiment divided over the number of transactions predicted by the initial and shifted FCEs. In every experiment, the number of transactions exceeds the number predicted by theory. On average the number of transactions is 30% more than predicted by the FCE volume. □

The efficiency phenomenon appears to be a direct result of speculative trading but as of yet, there is no clear theory that explains why so much speculative trading occurs. Nevertheless, speculation is a potential explanation for why complex markets such as these can achieve the high levels of efficiency seen in Table 2.

Price Levels

The relationships among trade prices and the two equilibrium concepts TE and FCE suggested in the Figures above can be made precise. Result 3 says that trade prices are not a constant as suggested by the competitive model that we apply. Result 4 initiates the analysis by telling us that prices are close to the two equilibria. Results 5 and 6 suggest how the two equilibrium concepts interact and how that interaction finds itself over time.

Result 3: The law of one price does not emerge under conditions of a constant FCE price.

Support: Without a formal definition of the law of one price we only refer to the figures that contain the time series of price data. The next result provides precise data.□

The close relationship between the models and the data suggested by Figures 4a-8b do indeed reveal themselves in the data.¹¹ Table 3 shows that the correlations between FCE price, TE price and trade price all exceed 0.96. While these correlations do reflect the influence of the basic principles as captured by the models, they also reflect the variation in FCE prices across experiments and price shifts and can thus be influenced by the choice of parameter shifts in an experiment. The contemporary relationships between price, TE price and FCE price that exit during periods of constant FCE price are better illustrated by the following results:

Table 3: Estimated Contemporaneous Cross-Correlations

	Trade Price	FCE Price	TE Price
Trade Price	1.00	0.96	0.98
FCE Price	0.96	1.00	0.98
TE Price	0.98	0.98	1.00

Result 4: (i) Traded prices are distributed around both FCE and TE prices. (ii) When trade prices deviate from the FCE price, they tend to deviate in the direction of the TE price.

¹¹ If the theoretical “consensus price” used in Goettler, Parlour, and Rajan (2005) is interpreted as the FCE price the following results indicate the extent to which observed market prices can be used as an estimate of the consensus price.

Support (i): The relationships among trade prices, FCE and TE can be illustrated with Figures 9, 10, and 11, which also provide general impressions of the data. Figure 9 shows the marginal distribution of trade prices around the FCE. Figure 10 shows the marginal distribution of trade prices around the TE. Figure 11 shows the marginal distribution of deviations in the TE from the FCE. The point of economic importance to learn from this relationship is that both FCE and TE appear to be “good” concepts of equilibrium.

Similarities exist among the distributions in Figures 9 and 10. Notice that the trade prices have “fat tails”. Trade prices appear to be T-distributed around the FCE and the TE. There is a statistically significant tendency for goods to be on average under priced relative to both the FCE and TE prices. Simple t-tests reject the null hypothesis that the mean of trade prices is equal to the FCE price at virtually any confidence level, but the economic significance, as well as the size of the under pricing in dollar terms is slight. Given a typical exchange rate of 500 francs (the currency of the experiment) = \$1, a 15-20 franc price deviation represents only about 3-4 cents and could easily be accounted for by subjective transaction costs.

Turning to Figure 11, the distribution of TE prices around the FCE, has properties similar to the distribution of trade prices around the FCE. However, trade prices have a higher variance than TE prices. TE prices have an estimated variance of 3654.7, while the estimated variance of trade prices is 8997.7, well over twice as high. The nature of this property is explored more closely by Result 5.

Figure 9: Distribution of Trade Prices Around FCE Price

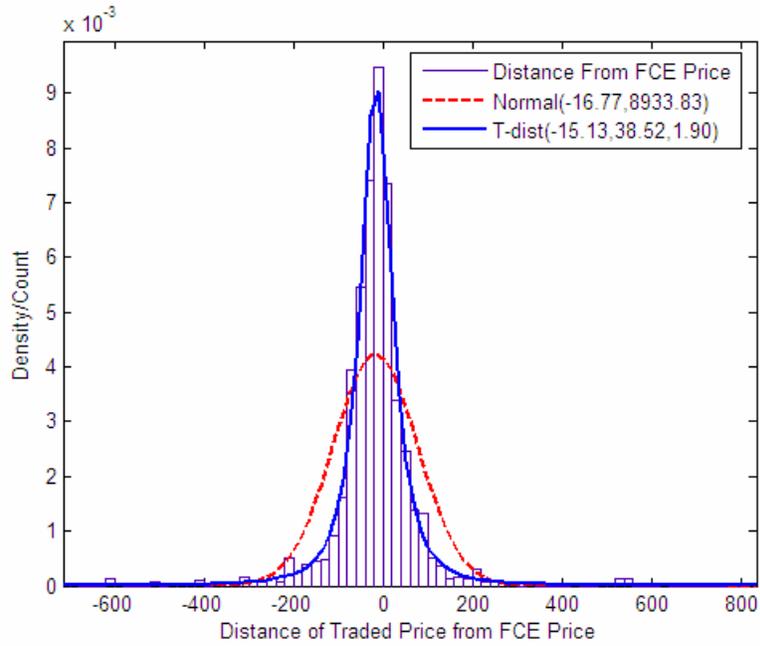


Figure 10: Distribution of Trade Prices Around the TE Price

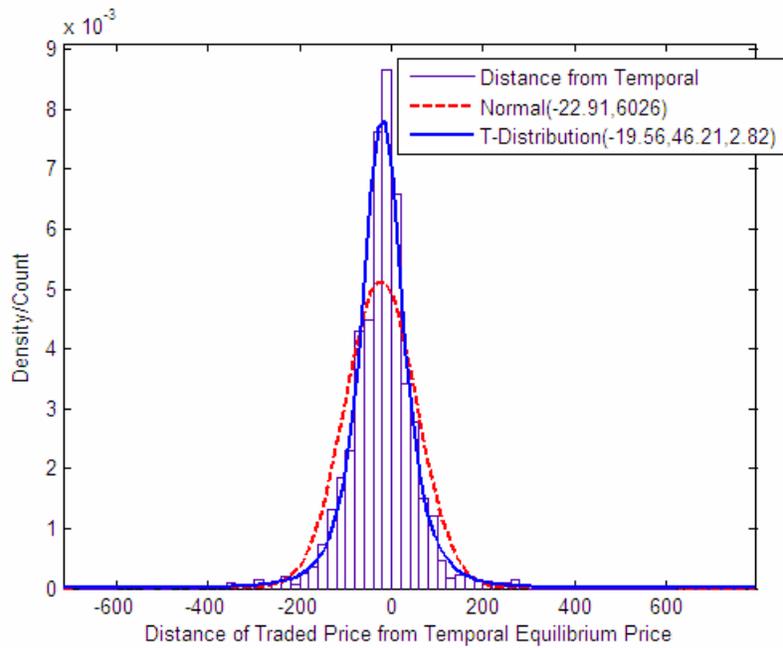
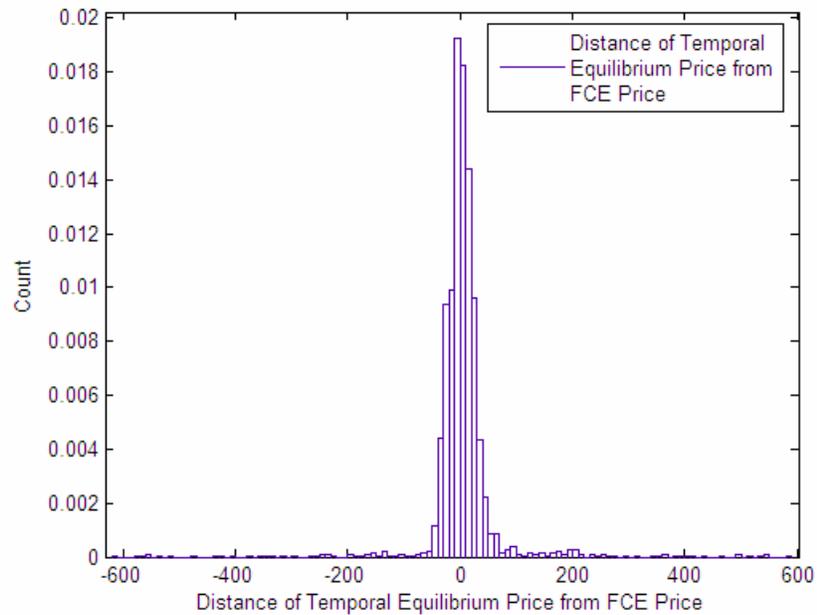
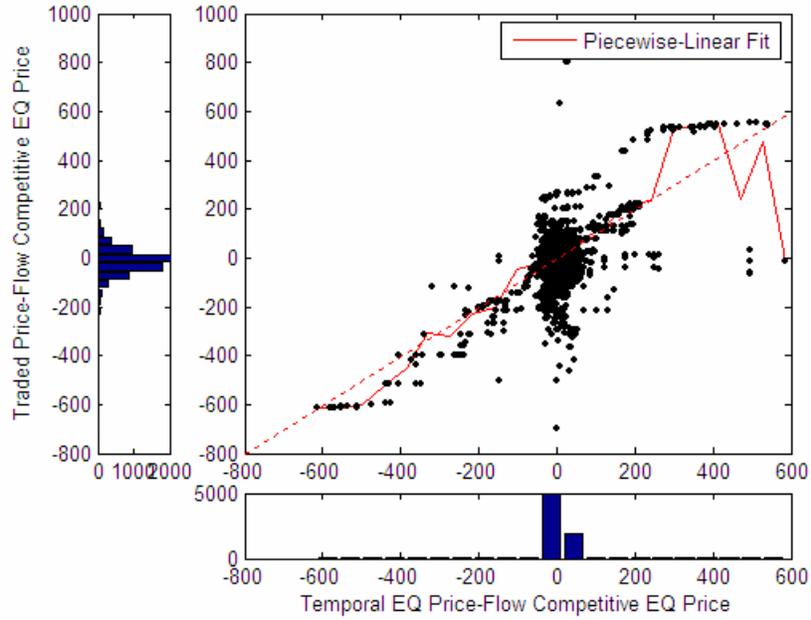


Figure 11: Distribution of TE Prices Around the FCE Price



Support (ii): Figure 12 illustrates the positive relationship between trade price deviations from the FCE and TE price deviations from the FCE. Across all experiments the contemporaneous correlation between these deviations is 0.6167. Notice that the relationship between temporal deviations and trade price deviations is weak when the TE is close to the FCE. This relationship becomes stronger when the TE deviations from the FCE are large in either direction.

Figure 12: Scatter Plot of Trade Price Deviations vs. TE Price Deviations from FCE



The following result spells out the relationship implicit in the discussion above. There is a complex interaction between the two equilibria. At the moment we can say only that it exists but of course the data beg for a theory to provide a description of the relationship.

Result 5: Both the direction of temporal equilibrium prices and the direction of the FCE price influence price movement.

Support: We use a simple least squares regression to predict future price movement based on how far away the current price is away from both the long run and the temporal equilibrium price for five different forecast horizons. Using only sections of data over which the FCE remains constant, we estimate the model:

$$P_{t+j} - P_t = \beta_0 + \beta_1(TE_t - P_t) + \beta_2(FCE - P_t) + \varepsilon_t$$

Where t, indexes the trade number.

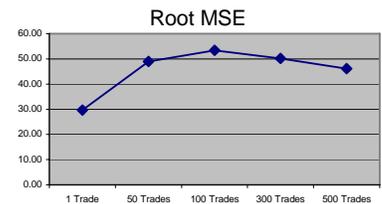
In this model, a slope coefficient of one is interpretable as “complete adjustment,” while a slope coefficient between zero and one indicates that prices are moving toward the equilibrium price, although not equilibrating perfectly.

Table 3 shows the results of these regressions for price changes after 1, 50, 100, 300, and 500 trades. The results indicate that, at all times, prices appear to be moving in the direction of both equilibriums since all of the estimated coefficients are between zero and one. The magnitude of these coefficients tends to grow with the forecast horizon, suggesting that prices, at least in the short run, are “sticky” and tend to under adjust over short time periods.

A different story emerges with the examination of price changes over much longer periods of time, 300 and 500 trades in the future. At these forecast horizons, the coefficients on the distance to the temporal equilibrium price and the distance to the FCE price sum to one, but both coefficients are statistically different from one. Neither equilibrium concept appears to dominate the other. Rather, each equilibrium appears to have its own distinct pull on prices. This property is undoubtedly a feature of some, more complicated equilibrium concept, which has yet to be theorized about. □

Table 4: FCE and TE in Forecasting Price Movement

Dependent Variable	Explanatory Variables			R Squared	Root MSE
	Constant	TE Price-Current Price	FCE Price-Current Price		
Price Change After:					
1 Trade	-2.28 (0.38)	0.05 (0.02)	0.08 (0.02)	0.07	29.72
50 Trades	-9.7 (0.65)	0.34 (0.03)	0.19 (0.03)	0.31	48.97
100 Trades	-10.93 (0.75)	0.16 (0.03)	0.49 (0.03)	0.36	53.43
300 Trades	-13.09 (0.91)	0.14 (0.04)	0.87 (0.04)	0.62	50.21
500 Trades	-18.88 (1.68)	0.38 (0.06)	0.61 (0.05)	0.72	46.14



If the price process in the market merely wandered randomly, one would expect the mean squared forecast error to grow unboundedly with the length of the forecast as

the process continued to accumulate random shocks. On the other hand, if the price were to follow an ARMA or other mean reverting process, we would expect the MSE forecast error to rise quickly at first and then level off to some constant. What we instead observe is the following result, which suggests a general equilibration property:

Result 6: The mean squared error of long-range price change forecasts based on the distance of current trade prices from the temporal equilibrium price and the FCE price is smaller than the mean squared error for short forecast horizons.

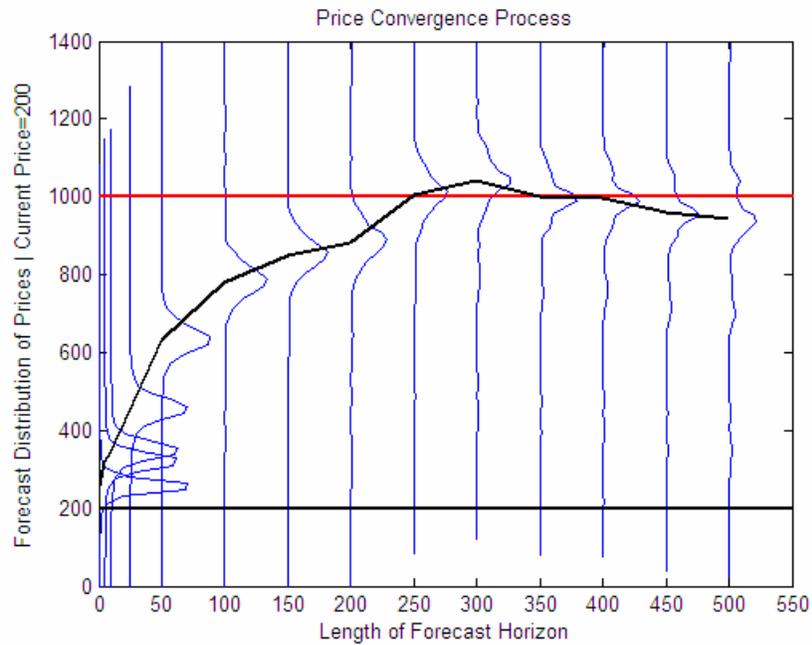
Support: What we see in Table 4 is that the mean squared error for the 500 trade-ahead price change forecast is actually smaller than the mean squared error for the 300, 100, and 50 trade-ahead forecasts. This suggests there is a “long range” equilibration, in which variation due to price stickiness or temporary perturbations is washed out.

Figure 13 illustrates the results of a simulation based on the results of Table 4 and the residual errors of the forecast equations. The coefficients in each equation of Table 4 specify the estimated mean of a price process, while the estimated matrix of residuals can provide an approximation of the distribution of prices around the mean. Using the estimated coefficients and the empirical distributions of the residuals, we can create price convergence surfaces.

The price convergence surface in Figure 13 below illustrates the probabilistic path of a good whose FCE and temporal equilibrium price at time t are $1000f$, and whose current price is arbitrarily set to $200f$. Each time- t cross section of the price convergence surface provides the estimated distribution of the traded price at time t .

By looking at Figure 13, we can see the results of Table 4 graphically, namely the probabilistic convergence of the price path and the evolution of the forecast error, increasing at first, then shrinking. □

Figure 13: The Price Convergence Surface



The structures of the price changes reported above lead naturally to issue of the role of bids and asks and the general dynamics of the price formation and price discovery process. These issues are explored in Alton and Plott (2007).

7. CONCLUDING REMARKS

The general conclusions of the experiments conducted here are very optimistic. First, it is clear that experimental methods in economics are capable of dealing with complex, continuous flow economies. Not only does the technology exist, experimental procedures are also satisfactory.

Second, complex markets with variable flows of incentives do not behave in a chaotic or unpredictable manner. The most basic question posed is whether there are tractable principles of economics that apply to changing flow environments. The answer is “yes”. On a local level, traded prices appear to track and respond to temporal equilibrium prices. We can say that price levels are well captured by supply and demand as defined by the temporal model. One could also say that there is a second law of supply and demand acting on markets which is defined by the probabilistic flow of incentives. Latent incentives, together with the arrival rates, produce flow concepts of demand and supply and an associated flow concept of a competitive equilibrium, which is the Flow Competitive Equilibrium (FCE). The FCE model has particularly strong predictive power for forecasting prices many trades into the future. At a deeper level however, these two laws appear act simultaneously and independently, each having a distinct “pull” on traded prices.

That the FCE model has such a pull, even after accounting for the influence of temporal order flow suggests a role of learning or speculation on the part of agents. This is supported by the fact that large numbers of speculative trades occur in each of the experiments conducted, with the number of trades achieved far exceeding the number predicted by incentive parameters alone.

Third, perhaps the most amazing feature of the experimental results reported above, is the tendency for these markets to achieve high levels of efficiency. A reasonable suspicion is that this is due primarily to speculation and market making smoothing prices over time, in a sense making the market less sensitive to short run order

flow imbalances, which would bring traded prices away from the FCE under conditions of no speculation or market making.

REFERENCES

- Aliprantis, R., D. Charalambos, and C. R. Plott, 1992, "Competitive Equilibria in Overlapping Generations Experiments", *Economic Theory*, 2:389-426.
- Alton, M. R. and C. R. Plott, 2007, "The Dynamics of Prices Adjustment in Experimental Markets with Continuous and Random Arrivals and Departures", California Institute of Technology, Pasadena, CA.
- Asparouhova, E., P. Bossaerts, and C. R. Plott, 2003, "Excess Demand and Equilibration in Multi-security Financial Markets: The Empirical Evidence." *Journal of Financial Markets* 6:1-21.
- Brewer P.J., M. Jiamg. B. Nelson, and C. R. Plott, 2002, "On the Behavioral Foundations of the Law of Supply and Demand: Human Convergence and Robot Randomness", *Experimental Economics*, 5:179-208.
- Cason, T.N. and D. Friedman, 1993, "Empirical Analysis of Price Formation in Double Auction Markets." *The Double Auction Market: Institutions, Theory and Evidence*, D. Friedman J. Rust. *SFI studies in Sciences of Complexity*. Redwood City, CA. Addison-Wesley.
- Cliff, D. and C. Preist, 1998, "Days without end: On the stability of experimental single-period continuous double auction markets." Hewlett-Packard Laboratories.
- Chamberlin, E.H., 1948, "An Experimental Imperfect Market", *Journal of Political Economy*, 56(2):95-108.
- Easley, D. and J. Ledyard, 1993, "Theories of Price Formation and Exchange in Oral Auctions," *The Double Auction Market: Institutions, Theory and Evidence*, D. Friedman J. Rust. *SFI studies in Sciences of Complexity*. Redwood City, CA. Addison-Wesley.
- Forsythe, R.E., T.R. Palfrey, and C.R. Plott, 1982, "Asset Valuation in an Experimental Market", *Econometrica*, 50:537-67.
- Friedman, D., 1991, "A Simple Testable Model of Double Auction Markets." *Journal of Economic Behavior and Organizations*, 16:47-70.
- Gode, D. and S. Sunder, 1993, "Allocative Efficiency of markets with zero intelligence traders: markets as a partial substitute for rationality." *The Journal of Political Economy* 101:119-137.
- Goettler, R., C. A. Parlour, and U. Uajan, 2005, "Equilibrium in a Dynamic Limit Order Market", *Journal of Finance*, LX, 5:2149- 2192.

- Jamison, J.C. and C. R. Plott, 1997, "Costly Offers and the Equilibration Properties of the Multiple unit Double Auction Under Conditions of Unpredictable Shifts of Demand and Supply", *Journal of Economic Behavior and Organization*, 31(4):591-612.
- Kagel, J.H., 2004, "Double Auction Markets with Stochastic Supply and Demand Schedules: Call Markets and Continuous Auction Trading Mechanisms." *Advances in Understanding Strategic Behavior: Essays in Honor of Werner Guth*, S. Hack, editor, Palgrave.
- Lian, P. and C.R. Plott, 1998, "General Equilibrium, Markets, Macroeconomics and Money in a Laboratory Experimental Environment". *Economic Theory* 12(1):21-75.
- Marimon, R. and S. Sunder, 1993, "Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence," *Econometrica*, 61(5):1073-1108.
- Miller, Plott and Smith 1977, "Intertemporal Competitive Equilibrium: An Empirical Study of Speculation", *Quarterly Journal of Economics* 91:599-624.
- Millner, E.L., M. D. Pratt and R. J Reilly, 1990, "Contestability in Real-Time Experimental Flow Markets", *Rand Journal of Economics*, 21(4):584-499.
- Parlour, C. A. 1998, "Price Dynamics in Limit Order Markets", *Review of Financial Studies*, 11:789-816.
- Peng L. and C. R. Plott, 1998, "General Equilibrium, Markets, Macroeconomics and money in a Laboratory Experimental Environment", *Economic Theory*, 12(1):21-75.
- Plott, C.R. and V. Smith 1978, "An Experimental Examination of Two Exchange Institutions", *The Review of Economic Studies*, XLV(1):133-53.
- C. R. Plott and P. Gray, 1990, "The Multiple unit Double Auction ", *Journal of Economic Behavior and Organization*, 13:245-58.
- Satterthwaite, M. and S. Williams, 1993, "The Bayesian Theory of the k-Double Auction" *The Double Auction Market: Institutions, Theory and Evidence*, D. Friedman J. Rust. *SFI studies in Sciences of Complexity*. Redwood City, CA. Addison-Wesley.
- Smith, V., 1962, "An Experimental Study of Competitive Market Behavior", *Journal of Political Economy*, 70(2):111-137.
- Smith, V., G.L. Suchanek, and A.W. Williams, 1988, "Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets, *Econometrica*, 56:1119-51.

Smith, V. and A. Williams, 1983, "An Experimental Comparison of Alternative Rules for Competitive Market Exchange" *Auctions, Bidding, and Contracting: Uses and Theory*. New York: New York University Press, 307-334.

Wilson, R., 1987, "On Equilibria of Bid-Ask Markets." *Arrow and the Ascent of Modern Economic Theory*. Ed. G. Feiwel, 375-414. UK: MacMillan.