

Clock Games: Theory and Experiments*

Markus K. Brunnermeier[†]
Department of Economics
Princeton University

John Morgan[‡]
Haas School of Business
and Department of Economics
University of California, Berkeley

This version: August 22, 2008

Abstract

In many situations, timing is crucial—individuals face a trade-off between gains from waiting versus the risk of being preempted. To examine this, we offer a model of clock games, which we then test experimentally. Each player’s clock starts upon receiving a signal about a payoff-relevant state variable. Since the timing of the signals is random, clocks are de-synchronized. A player must decide how long, if at all, to delay his move after receiving the signal. We show that (i) delay decreases as clocks become more synchronized and (ii) when moves are observable, players “herd” immediately after any player makes a move. Our experimental results are broadly consistent with these two key predictions of the theory.

Keywords: Clock games, experiments, currency attacks, bubbles, political revolution

*We thank Dilip Abreu, José Scheinkman and especially Thorsten Hens for helpful comments. Wiola Dziuda, Peter Lin, Jack Tang, and Jialin Yu deserve special mention for excellent research assistance. We also benefited from seminar participants at the University of Pittsburg, the Institute of Advanced Studies at Princeton, the UC Santa Cruz, the University of Michigan, UC Berkeley, and UCLA as well as from conference presentation at Princeton’s PLESS conference on Experimental Economics. The authors acknowledge financial support from the National Science Foundation. The first author thanks the Alfred P. Sloan Foundation for financial support while the second author thanks Trinity College, Cambridge for their generous hospitality.

[†]Department of Economics, Bendheim Center for Finance, Princeton University, Princeton, NJ 08544-1021, e-mail: markus@princeton.edu, <http://www.princeton.edu/~markus>

[‡]Haas School of Business, Berkeley, CA 94720-1900, e-mail: morgan@haas.berkeley.edu, <http://faculty.haas.berkeley.edu/rjmorgan/>

1 Introduction

It is often said that to succeed in business and in life, timing is crucial. Certainly this is true in terms of new product launches, stock picking, and real estate development. It is also true in choosing a mate or even starting a revolution. These situations all exhibit a trade-off between gains from waiting versus the risk of being preempted. In this paper, we offer a theoretical model of “clock games” that tries to capture this trade-off. We then test the key implications of the model in the lab.

In a clock game, each player’s clock starts at a random point in time: when he receives a signal of a payoff-relevant state variable (e.g., an opportune time for a product launch). Owing to this randomness, players’ clocks are de-synchronized. Thus, a player’s strategy crucially hinges on predicting the timing of the other players’ moves—i.e., predicting other players’ clock time. The exact prediction depends crucially on the observability of moves, the speed of information diffusion, and the number and size of players. There are n players making moves in our model; thus, our setting is the oligopoly analog to the “competitive” model of Abreu and Brunnermeier (2003), where there are continuum of small players.

While clock games are potentially quite complex to analyze, we transform the problem to one that can be readily analyzed using auction theory. In particular, the equilibrium waiting time in a clock game is isomorphic to the equilibrium bid in a multi-unit reverse first price auction with a stochastic outside option.

When moves are unobservable, the unique symmetric equilibrium in a clock game is remarkably simple—each player waits a fixed amount of time after receiving a signal before making a move. Slower information diffusion leads to longer equilibrium waiting time. When moves are observable, equilibrium waiting still has the same properties up to the time the first player moves. However, following this, herding occurs—all remaining players make their move immediately.

To test the theory, we run a series of experiments designed to examine the behavioral validity of two key synchronization factors: the speed of information diffusion and the observability of moves. To the best of our knowledge, we are the first to study these questions using controlled experiments.

The main results of the experiments are:

1. Equilibrium delay is robust—we observe delay in all treatments.
2. When moves are observable, there is initial delay followed by herding.
3. The slower the information diffusion, the longer the observed delay.

The remainder of the paper proceeds as follows: The rest of this section places clock games in the context of the broader literature on timing games. Section 2 presents the model, characterizes equilibrium play, and identifies key testable implications. Section

3 outlines the experimental design, while Section 4 presents the results. Finally, Section 5 concludes. Proofs of propositions as well as the instructions given to subjects in the experiment are contained in the appendices.

Related Literature At a broad level, clock games are a type of timing game (as defined in Osborne (2003)). As pointed out by Fudenberg and Tirole (1991), one can essentially think about the two main branches of timing games—preemption games and wars of attrition—as the same game but with opposite payoff structures. In a preemption game, the first to move claims the highest level of reward, whereas in a war of attrition, the last to move claims the highest level of reward.

Preemption games have been prominently used to analyze R&D races (see, e.g. Reinganum (1981), Fudenberg and Tirole (1985), Harris and Vickers (1985) and Riodan (1992)). In addition, a much-studied class of preemption games is the centipede game, introduced by Rosenthal (1981). This game has long been of interest experimentally, as it illustrates the behavioral failure of 1s induction (see e.g. McKelvey and Palfrey (1992)). In clock games (with unobservable moves), private information (which leads to the de-synchronization of the clocks) plays a key role, whereas centipede games typically assume complete information.¹ Indeed, this informational difference is crucial in the role that backwards induction plays in the two games. Since there is no *commonly known* point from which one could start the backwards induction argument, this rationale does not appear in clock games.

Clock games are also related to wars of attrition, where private information features more prominently. Surprisingly, there has been little experimental work on wars of attrition; thus, one contribution of our paper is to study the behavioral relevance of private information in a related class of games. Perhaps the most general treatment of this class of games is due to Bulow and Klemperer (1999), who generalize the simple war of attrition game by viewing it as an all-pay auction. Viewed in this light, our paper is also somewhat related to costly lobbying games, see e.g. Baye, Kovenock, and de Vries (1993), and the famous “grab the dollar” game, see e.g. Shubik (1971), O’Neill (1986), and Leininger (1989). Finally, the herding behavior, which is present in the clock games model with observable moves, is a feature also shared by Zhang (1997), whose model can be viewed as a war of attrition.

A recent paper by Park and Smith (2003) bridges the gap between these two polar cases by considering intermediate cases where the K th to move claims the highest level of reward.² The payoff structure of our clock game is as in Park and Smith: rewards are increasing up to the K th person to move and decreasing (discontinuously in our case) thereafter. In contrast to Park and Smith, who primarily focus on complete

¹An important exception is Hopenhayn and Squintani (2006), who study preemption games in an R&D context where each firm’s technological progress is stochastic and privately known.

²See also Park and Smith (2004) for leading economic applications of this model.

information, our concerns center on the role of private information and, in particular, on how private information results in de-synchronized clocks.

The key strategic tension in clock games—the timing of other players’ moves—figures strongly in the growing and important literature modeling currency attacks. Unlike clock games, which are inherently dynamic, the recent currency attack literature has focused on static games. Second generation models of self-fulfilling currency attacks were introduced by Obstfeld (1996). An important line of this literature begins with Morris and Shin (1998), who use Carlsson and van Damme’s (1993) global games technique to derive a unique threshold equilibrium. The nearest paper in this line to clock games is Morris (1995), who translates the global games approach to study coordination in a dynamic setting. The approach of Morris and Shin (1998) has spawned a host of successors using similar techniques as well as a number of experimental treatments (see, for instance, Heinemann, Nagel, and Ockenfels (2004) and Cabrales, Nagel, and Armenter (2002)).

As was described above, the clock games model is the oligopoly analog to the models in Abreu and Brunnermeier (AB 2002, 2003), who study persistence of mispricing in financial markets with a continuum of informationally small, anonymous traders.

2 Theory

2.1 Model

There are I players in the game. At any moment in (continuous) time, a player can decide when to exit. The game ends once $K < I$ players have exited.³ If a player exits at time t , he receives an “exit” payoff, e^{gt} . If, on the other hand, a player does not exit and the game ends, then he receives an “end of game” payoff, e^{gt_0} . The end of game payoff is stochastic as t_0 is exponentially distributed with parameter λ . Once a player exits, he cannot subsequently return; thus, each player’s strategy amounts to a simple stopping time problem. Tension in the model occurs for exit decisions at time $t > t_0$. By waiting, a player’s exit payoff increases at rate g , but risks the possibility of a payoff drop (from e^{gt} to e^{gt_0}) should the game end suddenly.⁴

In making his exit decision, each player receives a conditionally independent private signal about the realization of t_0 . Specifically, at time t_i , drawn uniformly from $[t_0, t_0 + \eta]$, player i learns that t_0 has already occurred. The parameter η , which is common knowledge, captures the speed of information diffusion that the t_0 event has

³If several players exit at the same instant in time, such that the number of those exiting exceeds K , we use a pro-rata tie-breaking rule to randomly determine which of the players exited successfully.

⁴To eliminate a “nuisance” equilibrium where all players wait forever before exiting, we assume that the game also ends at latest at time $t_0 + \bar{\tau}$, where $\bar{\tau}$ is a commonly known (large) constant. We assume that $\bar{\tau}$ is sufficiently large that the game always ends endogenously.

occurred. The interval $[t_0, t_0 + \eta]$ can be thought of as the “window of awareness” since all players receive signals during this timeframe.

Notice that player i can infer that t_0 lies between $t_i - \eta$ and t_i . Moreover, while player i does not know (the timing of) other players’ signals precisely, he can infer that by time $t_i + \eta$ all other players must have received their signals as well. Figure 1 illustrates the relationship between t_0 , signals, and payoffs. In this figure, t_0 has occurred at $t = 130$. The payoff growth rate, g , is 2%. Notice for $t < t_0$, the end of game payoff (stochastic dotted line) exceeds the exit payoff (solid curve) while for $t > t_0$ the reverse is true. Since $\eta = 50$ in the figure, the window of awareness is illustrated by the gray rectangle—all players receive private signals at $t \in [130, 180]$.




Figure 1: ‘Exit payoff’ versus ‘end of game payoff’

For the model to be interesting, the following assumptions are sufficient: (i) $0 < \lambda < g$, and (ii) η not too large. Assumption (i) guarantees that there is sufficient upside to waiting, and so strategic delay becomes a possibility. Assumption (ii) is needed to prevent the possible lag in the time a player receives a signal from becoming too large. Were this assumption violated, then the risk of a drop in payoff prior to receiving a signal would be sufficiently large that players would always choose to exit prior to receiving the signal. This assumption may be stated more precisely as follows: Let $\bar{\eta}$ solve $\Phi(K, I, \bar{\eta}\lambda) = \frac{Ig}{Ig - (I - K + 1)\lambda}$, where the function $\Phi(a, b, z)$ is a Kummer

hypergeometric function (see e.g. Slater (1974)).⁵ From the monotonicity properties of $\Phi(\cdot)$, such a solution always exists and is unique. Assumption (ii) requires that $0 < \eta < \bar{\eta}$.

Next, we characterize symmetric perfect Bayesian equilibria for two cases of the model. In the unobservable actions case, the only information a player has is her signal. In the observable actions case, in addition to her signal, each player learns of the exit of any other player. Formally, if player i exits at time t , then all other players observe this event at time $\lim_{\delta \rightarrow 0} (t + \delta)$.

2.2 Unobservable Actions

While one can solve for equilibria in clock games directly, a more elegant treatment recognizes that the game with unobservable actions can be recast as a static auction where players submit bid schedules as a function of their types ex-ante. This is similar to the approach often used in analyzing wars of attrition (see e.g. Fudenberg and Tirole (1991)).

Consider the following auction: Each player has a type t_i which is informative about the outside option e^{gt_0} with the same signal generating process as described above. Players simultaneously submit bids subject to the restriction that player i 's bid must weakly exceed e^{gt_i} . The K lowest bidders receive their bid amounts while the remainder receive the value of the outside option. Hence, this auction is a reverse first-price auction with multiple goods and a common value outside option.

Suppose that all other players use a bidding strategy $e^{g\beta(t_i)}$. Letting F_K be the cdf of the K th order statistic from $I - 1$ draws and f_K its associated density, then the expected payoff to bidder i when he bids e^{gb_i} is his bid when the K th player receives his signal after $\beta^{-1}(b_i)$ and the expected outside option otherwise:

$$E\pi_i = \int_{\beta^{-1}(b_i)}^{\infty} e^{gb_i} f_K(t_K|t_i) dt_K + \int_{-\infty}^{\beta^{-1}(b_i)} E[e^{gt_0}|t_i, t_K] f_K(t_K|t_i) dt_K.$$

Differentiating with respect to b_i , we have

$$[1 - F_K(\beta^{-1}(b_i)|t_i)] \beta'(\beta^{-1}(b_i)) g e^{gb_i} = f_K(\beta^{-1}(b_i)|t_i) \{e^{gb_i} - E[e^{gt_0}|t_i, t_K = \beta^{-1}(b_i)]\}. \quad (1)$$

Equation (1) reflects the following trade-off. An incremental increase in player i 's bid (in terms of waiting time) produces a marginal benefit of $\beta'(\beta^{-1}(b_i)) g e^{gb_i}$ provided

⁵Many of the solutions to the model involve integral terms of the form

$$\Phi(a, b, x) = \frac{(b-1)!}{(b-a-1)!(a-1)!} \int_0^1 e^{xz} z^{a-1} (1-z)^{b-a-1} dz.$$

In the appendix, we describe some useful properties of Kummer functions.

that he is among the K lowest bids. However, the risk associated with such an increase in waiting time is that player i could be tied for the K^{th} lowest bid. In that case, player i 's payoff drops from the inside option e^{gb_i} to the expected value of the common outside option, $E[e^{gt_0} | t_i, t_K = \beta^{-1}(b_i)]$.

Imposing symmetry, $b_i = \beta(t_i)$, we obtain the following ordinary differential equation

$$[1 - F_K(t_i | t_i)] \beta'(t_i) g e^{g\beta(t_i)} = f_K(t_i | t_i) \{e^{g\beta(t_i)} - E[e^{gt_0} | t_i, t_K = t_i]\}.$$

Now recall that $h_K(\cdot) := \frac{f_K(t_i | t_i)}{1 - F_K(t_i | t_i)}$ is the hazard rate for the K^{th} lowest of $I - 1$ draws. In the Appendix A.1 we establish that under our uniform-exponential information structure $h_K(\cdot)$ and $E[e^{-g(t_i - t_0)} | t_i, t_K = t_i]$ are independent of type t_i . Hence,

$$g\beta'(t_i) e^{g\beta(t_i)} - h_K e^{g\beta(t_i)} = -h_K (E[e^{-g(t_i - t_0)} | t_i, t_K = t_i] e^{gt_i}).$$

Solving this differential equation yields

$$\beta(t_i) = t_i + \frac{1}{g} \log \frac{h_K E[e^{-g(t_i - t_0)} | t_i, t_K = t_i] - \exp[(h_K - g)(t_i - C_1)]}{h_K - g},$$

where C_1 is a constant to be determined. Since $\beta(t_i) \geq t_i$ for all t_i , it then follows that the unique symmetric bidding equilibrium is where all bidders wait a fixed time, τ , after receiving their signals.⁶ That is

$$\beta(t_i) = t_i + \tau,$$

with

$$\tau = \frac{1}{g} \left(\log \frac{h_K}{h_K - g} + \log E[e^{-g(t_i - t_0)} | t_i, t_K = t_i] \right). \quad (2)$$

As we show in the appendix, the expression $\log \frac{h_K}{h_K - g} + \log E[e^{-g(t_i - t_0)} | t_i, t_K = t_i]$ may be expressed as a ratio of Kummer functions. This formulation proves particularly useful for comparative statics analysis. Thus, we have shown:

Proposition 1 *In the unique symmetric equilibrium each player waits for time τ to elapse after receiving the signal and then exits, where*

$$\tau = \frac{1}{g} \log \left(\frac{\lambda \Phi(K, I, \eta(\lambda - g))}{g - (g - \lambda) \Phi(K, I, \eta\lambda)} \right). \quad (3)$$

⁶In principle, the game could admit another equilibrium which is symmetric but not strictly increasing. In this equilibrium, all players wait infinitely long before exiting. The assumption that the game ends latest at $t_0 + \bar{\tau}$, which is equivalent to the imposition of the secret reserve price that excludes bids in excess of $e^{g(t_0 + \bar{\tau})}$, rules this out.

While the derivation above shows that τ is locally incentive compatible, we rule out global deviations in a (tedious) proof contained in Brunnermeier and Morgan (2006). The main point of Proposition 1 is to show that equilibrium behavior entails each player delaying some fixed amount of time *after* receiving the signal before exiting.

How does equilibrium behavior change as the clocks become less synchronized? To answer this question, it is useful to examine the relationship between the equilibrium delay, τ , and the size of the window of awareness, η . Figure 2 depicts this relationship for the parameters we use in the experiment, $I = 6$, $K = 3$, $g = 2\%$, and $\lambda = 1\%$.




Figure 2: Equilibrium delay, τ , for different η

As the figure shows, equilibrium delay is increasing in the length of the window of awareness. The main intuition is that, by making it more difficult for a player to predict the time at which others received the signal, a longer window of awareness blunts the preemption motive. Specifically, if a player knows exactly the time at which others received signals, then that player’s best response is to “undercut” the would-be pivotal player by exiting an instant before that player. Mutual undercutting reduces equilibrium delay. However, as η increases, this exercise becomes increasingly difficult. Since the marginal benefit of waiting, g , does not vary with the window of awareness, the reduction in the value of preemption (or equivalently in the marginal cost of waiting) leads to greater equilibrium delay. While Figure 2 illustrates this effect for particular parameters of our model, the result holds more generally.

Proposition 2 *Equilibrium delay is increasing in the length of the window of awareness.*

The relationship highlighted in Proposition 2 is one of the two main hypotheses we test experimentally.

Indeed, as a consequence of the monotonicity properties of the Kummer function in Equation (3), it is straightforward to show the following additional comparative static properties:

1. Equilibrium delay is increasing in K , the number of “exit slots.”
2. Equilibrium delay is decreasing in I , the number of players.

The intuition for the above comparative statics is straightforward. When exit slots are relatively abundant, players “bid” less aggressively and equilibrium delay increases. The reverse is true when exit slots are relatively scarce.

A more subtle comparison occurs when we fix the relative scarcity of slots at $\kappa := K/I$ and scale the game proportionately. In the limit, as $I \rightarrow \infty$, equilibrium delay has a strikingly simple form:

$$\lim_{I \rightarrow \infty} \tau = \frac{1}{g} \log \left(\frac{\lambda e^{\kappa \eta (\lambda - g)}}{g - (g - \lambda) e^{\kappa \eta \lambda}} \right) \quad (4)$$

Equation (4) has two key features. First, it is *identical* to that obtained in the AB model.⁷ Second, equilibrium delay is shorter in any finite sized version of the game than it is in the limit.

Why is delay shorter in the limit game? While the marginal benefit of a fixed delay is unchanged when the game is scaled, the marginal cost of delay increases owing to an increase in the chance of the game ending at the next instant. This chance depends (negatively) on the gap between the $K - 1$ and K th order statistics of the other players’ types. As the number of other players gets larger, this gap shrinks.

2.3 Observable Actions

We saw above that in clock games where moves are unobservable, equilibrium behavior entails delaying a fixed amount of time after receiving the signal before exiting. However, in many situations of economic interest, players are able to observe each other’s actions. We now explore how observability affects strategic delay.

Analogous to the case where actions are unobservable, we study equilibria where, prior to the first exit, all players wait a constant number of periods, τ_1 . After observing the first exit, it becomes commonly known that the highest possible value of the outside option lies strictly below the current price. Thereafter the game is strategically

⁷The equivalence between the two models is not immediate since the limit clock game consists of a countable infinity of players while in AB there is a continuum of players.

equivalent to (multi-good) Bertrand competition—the lowest $K - 1$ “bidders” enjoy the payoff from the inside option while the remainder enjoy the inferior outside option payoff. As a result, the unique equilibrium is the equivalent of marginal cost pricing—all of the remaining players optimally exit immediately after observing the first exit. Lemma 1 formalizes this intuition.⁸

Lemma 1 *In any perfect Bayesian equilibrium where the first player exits τ_1 periods after receiving the signal, all other players exit immediately upon observing this event.*

The herding result offered in Lemma 1 is analogous to that which occurs in the “dirty faces” game (Littlewood (1953)). When the current price exceeds the type of the last informed player, each individual knows that the inside option is worth strictly more than the outside option. However, since no player has yet exited, a player is unsure whether other players know this as well. Once a player has exited, the fact becomes commonly known and the game is transformed into pure Bertrand competition. A key testable implication of Lemma 1 is that equilibrium behavior will necessarily give rise to herding following the decision of the first player to exit.

Of course, the model where exit is totally unobservable and the present situation, where exit is perfectly observable, represent the two extreme cases. Realistic situations will tend to lie somewhere between these two. Together, Proposition 1 and Lemma 1 suggest that the greater the observability of the exit decision, the more bunched are the exit times.

Next, we turn to the timing of the exit decision prior to the first exit. Prior to the first exit, the game is strategically equivalent to a low-bid auction for reasons identical to the unobservable case. Using Lemma 1, we can identify the expected payoff in the continuation game where player i is not the first to exit. Player i 's expected payoffs are then

$$\begin{aligned} E\pi_i &= \int_{\beta^{-1}(b_i)}^{\infty} e^{gb_i} f_1(t_1|t_i) dt_1 \\ &\quad + \frac{I-K}{I-1} \int_{-\infty}^{\beta^{-1}(b_i)} E[e^{gt_0}|t_i, t_1] f_1(t_1|t_i) dt_1 \\ &\quad + \frac{K-1}{I-1} \int_{-\infty}^{\beta^{-1}(b_i)} e^{g\beta(t_1)} f_1(t_1|t_i) dt_1. \end{aligned}$$

where t_1 is the lowest of $I - 1$ types and $\beta(\cdot)$ is a strictly increasing bidding strategy. Notice that the expected payoff reflects the fact that, following the first exit, each remaining player has an equal chance $\frac{K-1}{I-1}$ of receiving the inside option.

⁸Herding, in this instance, arises from the fact that the private information of the first player to exit is (partially) revealed by his decision to exit. This is analogous to the signaling role of the timing of moves which is prominent in Chamley and Gale (1994) as well as Gul and Lundholm (1995).

Differentiating with respect to b_i , we have

$$[1 - F_1(\beta^{-1}(b_i) | t_i)] \beta'(\beta^{-1}(b_i)) g e^{g b_i} = f_1(\beta^{-1}(b) | t_i) \frac{I - K}{I - 1} \{e^{g b_i} - E[e^{g t_0} | t_i, t_1 = \beta^{-1}(b_i)]\}.$$

Following steps identical to those when actions are unobservable, we obtain a unique symmetric equilibrium stopping time. In equilibrium, each player waits for time τ_1 after receiving their signal before exiting, where

$$\tau_1 = \frac{1}{g} \left[\log \frac{h_1}{h_1 - g \frac{I-1}{I-K}} + \log E[e^{-g(t_i - t_0)} | t_i, t_1 = t_i] \right]. \quad (5)$$

Comparing Equation (5) with Equation (2), the analogous expression when actions are unobservable, one notices two key differences: First, g is replaced by $\frac{I-1}{I-K}g$ in the first log-term in Equation (5). This reflects the fact that, even after the first player exits, all remaining players have a $(K - 1)$ to $(I - 1)$ chance of getting out at the high payoff in the next instant. Second, the hazard rate of a drop in payoff is equal to the conditional probability that the first player will exit in the next instant. In other words, the hazard rate is identical to above h_K if one sets $K = 1$. Finally, note that the term $E[e^{-g(t_i - t_0)} | t_i, t_1 = t_i]$ is the same for both settings. Using steps analogous to those leading to Proposition 1 allows us to derive τ_1 in closed form and thereby characterize a unique symmetric equilibrium to the game.

Proposition 3 *In the unique symmetric equilibrium, if no players have exited, each player waits for time $\tau_1 > 0$ to elapse after receiving the signal and then exits, where*

$$\tau_1 = \frac{1}{g} \log \left(\frac{\lambda \Phi(1, I, \eta(-g + \lambda))}{\frac{I g}{I - K + 1} - \left(\frac{I g}{I - K + 1} - \lambda \right) \Phi(1, I, \eta \lambda)} \right).$$

Once any player has exited, all other players exit immediately.

Proposition 3 has in common with Proposition 1 the feature that it is optimal for a player to delay exiting for a period of time after receiving the signal. Indeed, some properties associated with equilibrium comparative statics for the unobservable case continue to hold in the observable case. For instance, following the same steps as in the proof of Proposition 2, one can readily show that equilibrium delay, τ_1 , is increasing in the length of the window of awareness for the observable case as well.

It is also interesting to consider how delay changes with the scale of the game when moves are observable. Again fixing the relative scarcity of exit “slots” and taking limits, we obtain

$$\lim_{I \rightarrow \infty} \tau_1 = 0 \quad (6)$$

As Equation (6) shows, there is no equilibrium delay in the limit game. The reason is that the hazard of the game ending becomes unbounded in the limit. For a fixed delay,

the marginal benefit is unchanged as the game scales; however, the marginal cost of waiting becomes unbounded in the limit.

Equilibrium delay for the observable case is always shorter than the unobservable case in the limit game. For finite clock games, however, the comparison is ambiguous. To see this, fix the parameter values of the model at $I = 6$, $K = 3$, $g = 2\%$, and $\lambda = 1\%$. Numerical calculations show that $\tau_1 > \tau$ for $\eta < 59.8360$ and $\tau_1 < \tau$ for $\eta > 59.8361$. Thus, while strategic delay is common to both cases, there is no systematic ordering between τ_1 and τ .

3 Experimental Design and Procedures

The experiment sought to closely replicate the theoretical environment of clock games. The experiment consisted of 16 sessions conducted at the University of California, Berkeley during spring and fall 2003. Subjects were recruited from a distribution list comprised of undergraduate students from across the entire university, who had indicated a willingness to be paid volunteers in decision-making experiments. For this experiment subjects were sent an e-mail invitation promising to participate in a session lasting 60-90 minutes, for which they would earn an average of \$15/hour.

Twelve subjects participated in each session, and no subject appeared in more than one session. Throughout the session, no communication between subjects was permitted, and all choices and information were transmitted via computer terminals. At the beginning of a session, the subjects were seated at computer terminals and given a set of instructions, which were then read aloud by the experimenter. A copy of the instructions appears in Appendix B.

Owing to the complexity of the clock game environment, we framed the experiment as a situation in which subjects played the role of “traders” deciding on the timing of selling an asset and receiving a signal that the price of the asset has surpassed its fundamental value. Thus, the end-of-game payoff, in this setting, corresponds to the fundamental value of the asset. The exit payoff is simply the current price of the asset at the time a trader sold it. Of course, this design decision comes with both costs and benefits. The main benefit is to speed learning by subjects by making the game more immediately understandable. Since our main interest is in testing equilibrium comparative statics arising from the theory, convergence to some sort of stable behavior is essential. A secondary benefit is that understanding trading decisions in environments characterized by stock price “bubbles” is of inherent interest. The cost, of course, is that the particular frame we chose for the clock game may drive the results.

Each session consisted of 45 “rounds” or iterations of the game, all under the same treatment.⁹ Subjects were informed of this fact. At the beginning of each round, subjects were randomly assigned to one of two “markets” consisting of six traders

⁹Owing to networking problems, session 3 lasted only 35 rounds.

each, (i.e. $I = 6$). The job of a trader was to decide at what price to sell the asset. Subjects saw the current price of the asset. While we cannot directly replicate the continuous time assumption of the model in the laboratory setting, we tried to closely approximate it.¹⁰ The price of the asset began at 1 experimental currency unit (ECU) and increased by 2% for each “period”, (i.e. $g = 2\%$), where periods lasted about a half second each.¹¹ The computer determined when the “true value” of the asset stopped growing. In each period, there was a 1% chance of this event (i.e. $\lambda = 1\%$). At a random period after the true value of the asset had stopped growing (described in detail below), a subject also received a message that “the price of the asset is above its true value.” Finally, in Observable treatments (described in more detail below), traders were also informed each time some other seller sold his unit of the asset.

Once three or more traders in a market sold the asset, the round ended, (i.e. $K = 3$).¹² Following this, subjects learned their current and cumulative earnings as well as the prices at which all of the assets were sold. A subject’s earnings in a round were determined as follows: If the subject successfully sold the asset (i.e., was among the first three traders to sell), he received the price of the asset at the time he sold it. Otherwise, the subject earned an amount equal to the “true value” of the asset (end-of-game payoff). The parameter values used in the experiment are: $I = 6$, $K = 3$, $g = 2\%$, $\lambda = 1\%$, and $\bar{\tau} = 200$. Subjects received \$1 for each 50 ECUs earned, with fractions rounded up to the nearest quarter. Earnings averaged \$15.16, and each session lasted from 50 to 80 minutes.

Treatments

We examined how changes in both observability and the window of awareness impact the timing of exit. In the Baseline treatment, each subject learned that the true value of the asset had stopped growing with a lag that was uniformly (and independently) distributed from 1 to 90 periods (i.e. $\eta = 90$). In the Compressed treatment, we reduced the window of awareness, η , from 90 to 50 periods. This provides a direct test of Proposition 2. In the Observable treatment the window of awareness was the same as in Baseline; however, subjects were informed whenever a trader sold the asset. Comparing Baseline to Observable enables a direct comparison of Propositions 1 and 3. We ran six sessions each under the Baseline and Compressed treatments and four

¹⁰While the experimental design is, of necessity, run in discrete time, we computed numerically equilibrium waiting times for the discrete version of the model and verified that they converged to those for the continuous time case in the limit.

¹¹One may worry that, since periods lasted about a half second each, a subject may have had insufficient time to react to the event of a sell by another subject with only a one period delay. However, most studies of reaction time to light stimuli for college age individuals indicate a mean reaction time of approximately 0.19 seconds. See, for instance Welford (1980).

¹²In principle, a round could also end if fewer than three traders sold the asset and 200 price “ticks” had elapsed after the price of the asset exceeded its “true value.” This never occurred in any round of any session.

sessions under the Observable treatment, giving 16 sessions overall. A total of 192 subjects participated in these experiments.

Experimental Design Rationale

A key consideration in the experimental design was to minimize information “leakage” about trading behavior in the experiment. In pilot experiments, we found that subjects used auditory cues (i.e. mouse clicks) to infer trading behavior in unobservable treatments. To remedy this problem, our experimental design had subjects sell by *hovering* (instead of clicking) their mouse over the sell button.¹³

While this minimized information leakage, it occasionally led to selling “mistakes” when a subject’s mouse pointer inadvertently strayed into the sell box. Many of these mistakes are fairly obvious: sales would sometimes occur within the first few periods of a round. In the case of the Baseline and Observable treatments, we “cleaned” the data by eliminating observations where sales occurred within the first 10 periods after the start of the round. In the case of the Observable treatment, we dropped a round entirely when the first sale occurred within the first 10 periods.

While the decision a subject faced in each round of the game—when to sell the asset—is relatively simple, the price and information generating process are complicated. We expected that subjects would require several rounds of “learning by doing,” so we ran 45 repetitions of the game in a stationary environment. The extensive feedback given to subjects after each round was also designed to speed learning. In addition, whenever a subject sold the asset *below* its true value, he received a message indicating this fact. To get a sense of how well subjects understood the game at the end of a session, we asked each subject to fill out a post-experiment questionnaire where they were asked to describe their strategy. In the vast majority of instances, subjects described their strategies as waiting for the price of the asset to rise a certain amount after receiving the message that the asset was above its true value and then selling.¹⁴

We worried that, by running 45 iterations of the game, it would, in effect, become a repeated game. To reduce this risk, we randomly and anonymously rematched subjects after each round of the game and prohibited communication among subjects. While it is theoretically possible for subjects to coordinate on dynamic trading strategies, achieving the required coordination struck us as difficult. In examining the data, we looked for evidence of “collusive” strategies on the part of subjects. Such strategies might consist of delaying an excessively long time to sell after receiving the signal or coordinating on a particular price of the asset at which to sell regardless of signals

¹³We considered forcing subjects to click to advance to the next period in the game or playing loud music to minimize the perception of clicks. The first approach has the disadvantage of slowing down the game and removing the continuous time element to decisions. The second approach has the disadvantage of distracting the subjects and “priming” them to play a certain way depending on the tempo of the music.

¹⁴The formal empirical analysis makes no use of the answers given in the questionnaire.

received. We found no evidence of either type of behavior. Further, no subject mentioned coordinating or dynamic strategies in their responses to the post-experiment questionnaire.

There was considerable variability in subject choices in the early rounds of each session; however behavior was more stable in the last 25 rounds. Since we are primarily interested in the equilibrium performance of the model rather than in learning, we confine attention in the results section below to these rounds.¹⁵

4 Results

In this section we present the results of the laboratory experiment. We are mainly interested in the following measures of subject choices:

1. *Duration*: We measure the length, in periods from t_0 until the end of the game—that is, the period in which the third seller sold the asset. In the event that the game ended in a period prior to t_0 , we code Duration as zero.
2. *Delay*: We measure the length, in periods, of strategic delay by sellers. The variable *Delay* for seller i is the number of periods between the time he received the signal until the time he sold the asset. If i never sold the asset, then *Delay* is coded as missing. If i sold at or before the time he received the signal, *Delay* equals zero.¹⁶
3. *Gap*: We measure the gap, in periods, between the sale times of the i th and $i + 1$ th subjects selling the asset.

The first two measures, Duration and Delay, enable us to study the main implication of Propositions 1 and 3—namely that equilibrium behavior will lead traders to engage in strategic delay. Indeed, the Delay measure is the empirical counterpart to the τ and τ_1 predictions derived in the theory. Further, the main implication of Proposition 2 is that a reduction in the window of awareness reduces both Duration and Delay. Finally, the measure Gap seeks to capture the key behavioral prediction of Lemma 1—that observable trading information leads to “herding” on the part of sellers following the first sale.

Our delay variable suffers from a problem identical to that encountered in studying bidding in ascending auctions—we only observe a censored measure of delay for

¹⁵This is not to say that the nature of learning behavior in clock games is uninteresting per se. However, given our concerns with equilibrium comparative static predictions of the theory model, we feel that a careful study of subject learning in clock games is beyond the scope of the present paper.

¹⁶We also investigated an alternative coding scheme whereby a missing value was assigned for the Delay of sellers who sold but never received the signal. The results are qualitatively unaffected by this alternative. Details are available from the authors upon request.

subjects who did not sell prior to the game ending. We deal with this in two ways. Conservatively, we code the delay variable as missing when it is censored. We also use a Tobit specification, reported in Table 3 below, to try to recover censored delay values. The results are qualitatively similar under either approach.

Table 1 presents the predictions of the theory model for each of these performance measures. The Delay measure is simply a numerical evaluation of the strategies offered in Propositions 1 and 3. The Gap measure equals the distance between adjacent order statistics, $\eta/7$, for the unobservable case. The Gap measure is 1 for the observable case owing to discrete time. To obtain the duration measure, recall that the i th lowest of 6 draws from a uniform distribution of length η is $(i/7)\eta$. Thus, the duration measure is simply $(3/7)\eta + \tau$ in the unobservable case and $(1/7)\eta + \tau_1 + 1$ in the observable case, where again the additional “+1” stems from discrete time. As Table 1 shows, the expected Duration is predicted to be longest in the Baseline treatment and shortest in the Compressed and Observable treatments. Delay is predicted to be much shorter under the Compressed or Observable treatments compared to Baseline. The Gap measure illustrates a distinct difference between the Baseline and Compressed treatments and the Observable treatment.

Table 1: Theory Predictions

	Treatment		
	Baseline	Compressed	Observable
Duration	62	26	28
Delay*	23	5	14
Gap	13	7	1

* For the Observable treatment, Delay is only meaningful for the first seller.

4.1 Overview

Table 2 presents descriptive statistics treating each session as an independent observation.

Table 2: Descriptive Statistics (Periods 21-45)

	Treatment		
	Baseline	Compressed	Observable
Number of Sessions	6	6	4
Duration	43.31 (8.42)	26.48 (1.56)	32.30 (4.52)
Delay			
Seller 1	6.97 (2.30)	3.99 (1.08)	6.59 (1.30)
Seller 2	10.14 (3.44)	5.26 (1.49)	
Seller 3	12.31 (4.51)	6.72 (1.29)	
Gap			
Between 1st & 2nd seller	23.47 (3.67)	18.57 (5.58)	4.20 (0.97)
Between 2nd & 3rd seller	15.45 (2.46)	8.68 (0.76)	1.86 (0.08)

Standard deviations in parentheses

Consistent with the theory, Duration is longest in the Baseline treatment and shortest in the Compressed treatment. Delay also mirrors the comparative static predictions of the theory: it is longest in Baseline and shortest in Compressed. The Gap measure reflects the main effect of the Observable treatment—sellers after the first are strongly clustered in their sell time. Compared to the Baseline and Compressed treatments, gaps in the Observable treatment are much shorter. Indeed, the gap between the second and third seller is extremely close to the theoretical prediction.




Figure 3: Duration by Treatment

To see the variation in the data within each session, we present histograms of Duration. In Figure 3, the bar associated with zero indicates the fraction of cases in which no “bubble” formed—that is, three subjects exited before t_0 . This event, which in theory should *never occur*, happens over 10% of the time in the Baseline and Observable treatments, but less than 5% of the time in Compressed.

Taken together, there is considerable evidence for a number of treatment effects predicted by the theory; however, there is considerable variability in outcomes and some notable discrepancies between the theory predictions and observed behavior. In the rest of this section, we perform a variety of statistical tests to understand choice behavior and treatment effects in more detail.

4.2 Session-Level Analysis

In this section, we study treatment effects treating the session as the unit of observation. Obviously, this is a conservative approach to the data as it reduces the dataset to 16 observations. Throughout, we rely on two types of statistical tests to formally investigate treatment effects. The first test is a Wilcoxon Rank-Sum (or Mann-Whitney)

test of equality of unmatched pairs of observations. This is a non-parametric test which gives back a z -statistic which may be used in hypothesis testing. Our second test is a standard t -test under the assumption of unequal variances. This test has the advantage of familiarity, but the disadvantage of requiring additional distributional assumptions on the data to be valid. As we will show below, the conclusions drawn from the two tests rarely differ for our data.

Prediction 1. *Duration is longer in the Baseline than in either the Compressed or the Observable treatments.*

Support for Prediction 1.

We test the null hypothesis of no treatment effect against the one-sided alternative predicted by the theory. Comparing Compressed to Baseline, we obtain a z -statistic of 2.88 and a t -statistic of 4.81. Both reject the null hypothesis in favor of the alternative hypothesis at the 1% significance level. Comparing Observable to Baseline, we obtain a z -statistic of 1.92 and a t -statistic of 2.68. Both reject the null hypothesis in favor of the alternative hypothesis at the 5% significance level.

Prediction 2a *Delay is longer in the Baseline than in the Compressed treatment.*

Support for Prediction 2a.

Since Table 2 suggested that the first, second, and third sellers behave somewhat differently, we test the null hypothesis of no treatment effect against the one-sided alternative implied by the theory separately for each seller.

Seller 1.

Comparing Compressed to Baseline, we obtain a z -statistic of 2.08 and a t -statistic of 2.87. Both reject the null hypothesis in favor of the alternative hypothesis at the 5% significance level.

Seller 2.

Comparing Compressed to Baseline, we obtain a z -statistic of 2.08 and a t -statistic of 3.19. Both reject the null hypothesis in favor of the alternative hypothesis at the 5% significance level.

Seller 3.

Comparing Compressed to Baseline, we obtain a z -statistic of 2.08 and a t -statistic of 2.92. Both reject the null hypothesis in favor of the alternative hypothesis at the 5% significance level.

Taken together, these results provide strong support at the session level for Proposition 2.

Prediction 2b. *Delay is longer in the Baseline than in the Observable treatment.*

Lack of Support for Prediction 2b.

For the comparison to be meaningful, we restrict attention to the first seller (since the theoretically relevant comparison is between τ and τ_1). Comparing Observable to Baseline, we obtain a z -statistic of 0.43 and a t -statistic of 0.33. Neither test rejects the null hypothesis of no treatment effect (p-values of 0.67 and 0.75, respectively).

Prediction 3. *Gap is longer in the Baseline than in either the Observable or the Compressed treatments.*

Support for Prediction 3.

Again, based on Table 2, we distinguish the gap between the first and second sellers from the gap between the second and third sellers.

Sellers 1 and 2

Comparing Compressed to Baseline, we obtain a z -statistic of 1.76 and a t -statistic of 1.80. Both reject the null hypothesis of equal Gaps in favor of the alternative hypothesis predicted by the theory at the 10% significance level. Comparing Observable to Baseline, we obtain a z -statistic of 2.56 and a t -statistic of 12.24. Both reject the null hypothesis of no treatment effect in favor of the alternative hypothesis predicted by the theory at the 1% significance level.

Sellers 2 and 3

Comparing Compressed to Baseline, we obtain a z -statistic of 2.88 and a t -statistic of 6.44. Both reject the null hypothesis in favor of the alternative hypothesis at the 1% significance level. Comparing Observable to Baseline, we obtain a z -statistic of 2.56 and a t -statistic of 13.51. Both reject the null hypothesis in favor of the alternative hypothesis at the 1% significance level.

To summarize, many of the key comparative static predictions of the theory are supported by the data—even treating the session as the unit of observation.

4.3 Individual-Level Analysis

Censoring

When the game ends before a player exits, our Delay measure is censored. In the session level analysis, we dealt with this by omitting such observations entirely. This potentially biases our measure of delay downward. By treating each subject/round as the unit of observation, we may employ techniques that account for this censoring problem to derive unbiased estimates of delay.

Suppose that a subject was planning to exit (sell) in period T_i^* . If the game ends in period $T^{\text{end-of-game}} \leq T_i^*$, our observation of delay for this subject is right-censored. However, if we use the equilibrium structure of the game, we can recover a measure of delay for that subject. Recall that the theory predicts that a subject will wait a fixed number of periods, τ , after receiving the signal before selling. Allowing for a Normally distributed error term, $\varepsilon_{ir} \sim \mathcal{N}(0, \sigma_{\text{treatment}}^2)$, then a subject i 's exit time in round r of the game is $T_{ir}^* = t_{ir} + \tau + \varepsilon_{ir}$ where t_{ir} is the time the signal was received.¹⁷ Rewriting this expression, we have

$$DELAY_{ir}^* = T_{ir}^* - t_{ir} = \tau + \varepsilon_{ir},$$

¹⁷Of course, strictly speaking the error term cannot be normally distributed since the observed T_i^* is always non-negative. However, this should not be problematic since τ is sufficiently large.

where, conditional on t_{ir} , the variable $DELAY^*$ is Normally distributed with mean τ and variance $\sigma_{\text{treatment}}^2$. With this specification, we can use the Tobit procedure to obtain estimates of τ despite the right-censoring.¹⁸ Since Tobit estimates are known to be quite sensitive to assumptions on the error structure, we allow for the possibility that the error term is different across treatments and run separate estimations for each treatment. For the Baseline treatment, delay is estimated to be 17.277 periods (standard error = 5.31)—considerably longer than the session level statistics offered in Table 2. For the Compressed treatment, Tobit produces a delay estimate of 11.011 periods (standard error = 7.46), which is again larger than the session level statistics. These corrected estimates for our Delay measure do not alter our conclusions with respect to Prediction 2a. We can reject the hypothesis that the two coefficients are equal at the 1% significance level in favor of the one-sided alternative predicted by the theory.

Early Exit

A central prediction of the theory is that players only exit after receiving their signal or, in the case of the Observable treatment, after observing the time of the first exit. Yet, as Figure 3 highlights, in some cases, subjects sell the asset *prior* to receiving the signal.

To understand the factors predicting the decision to exit prior to receiving a signal, we performed a probit analysis where the dependent variable, *EARLY-EXIT*, is a dummy which equals 1 if a subject sold (weakly) prior to receiving the signal. Our regressors are dummies for each treatment, a linear term for the time the asset first exceeded its true value (t_0) and the interaction between these variables. It is difficult to interpret Probit coefficients directly; thus, Table 3 reports the results expressed in terms of the changes in the probability of early exit. Column 1 excludes the results of the Observable treatment since, in that treatment, the equilibrium calls for all traders to sell immediately after the first sale—regardless of whether they received the signal. Column 2 includes all treatments but restricts the sample to the first seller.

¹⁸Owing to differences in the information structure, the Observable treatment is not amenable to this analysis.

Table 3: Probit Model of Probability of *EARLY-EXIT*

	Baseline and Compressed Only	Seller 1 Only
Compressed	0.029* (0.0394)	0.145 (0.1065)
Observable		0.230* (0.1115)
t_0	0.003** (0.0003)	0.007** (0.0008)
$t_0 \times$ Compressed	-0.001** (0.0003)	-0.002* (0.0009)
$t_0 \times$ Observable		-0.003** (0.0010)
Baseline Probability	0.09	0.29
Round Fixed Effects	Yes	Yes
Observations	2259	738

Robust z -statistics in parentheses
 * significant at 5%; ** significant at 1%

To obtain estimates of the marginal effect of each treatment on the probability of early exit, one needs to add the coefficient estimate for the treatment with the estimate for the interaction term multiplied by the average value of t_0 , which is about 75. The column 1 estimates imply that the marginal effect (i.e., the interaction of time with session type) of the Compressed treatment is to *reduce* the probability of early exit. The column 2 estimates imply that, for first sellers, the Compressed and Observable treatments modestly reduce the probability of early exit compared to Baseline.

More surprising is the role that t_0 has on the probability of early exit. Regardless of the sample, the coefficients imply that the larger the value of t_0 , the greater the probability that a subject exited before receiving the signal. This is not predicted by the theory but does correspond with a heuristic “satisficing” strategy where subjects lock-in gains once the value of the asset has reached some aspiration level.

Delay

In Table 3, we saw that the timing of t_0 was a key driver leading players to sell before receiving the signal. Here, we investigate whether it plays a similar role for observations where we have an uncensored measure of delay. Specifically, for this subsample, we

regress delay on dummies for each treatment, a linear t_0 term interacted with treatment. We also add round fixed effects to control for learning. Finally, to deal with possible heteroskedasticity and autocorrelation, we report standard errors clustered by subject. The results of this analysis are shown in Table 4.

Table 4: Delay Estimates

	Seller 1	Seller 2	Seller 3
Constant	12.841** (10.78)	17.931** (10.30)	22.803** (10.28)
Compressed	-6.861** (5.16)	-10.281** (5.22)	-13.023** (5.17)
Observable	-2.169 (1.07)		
t_0	-0.071** (9.59)	-0.097** (8.17)	-0.127** (8.81)
$t_0 \times$ Compressed	0.045** (5.03)	0.064** (4.59)	0.086** (4.82)
$t_0 \times$ Observable	0.012 (0.93)		
Round Fixed Effects	Yes	Yes	Yes
Observations	738	584	583
R -squared	0.23	0.26	0.28

OLS: Robust t -statistics in parentheses.

* significant at 5%; ** significant at 1%

As the table shows, our findings regarding Predictions 2a and 2b continue to be borne out even after controlling for t_0 and learning effects. Delay is shorter under the Compressed treatment than Baseline as the theory predicts; however there is no statistical difference between Baseline and Observable.

The variable t_0 , which is theoretically irrelevant, does appear to influence subject choices. In particular, the coefficient estimates indicate that larger values of t_0 are associated with significantly less Delay in all treatments.¹⁹ Indeed, we can reject the null hypothesis of a zero t_0 effect against the one sided alternative at the 5% significance level for all specifications. The t_0 effect is, however, systematically less pronounced in the Compressed treatment compared to Baseline or Observable.

¹⁹To compute the economic magnitude of the an incremental change in t_0 , one needs to add the t_0 coefficient to the coefficient of the interaction term for the relevant treatment.

5 Conclusions

We introduced and analyzed a class of games we refer to as clock games. Players in these games must decide the timing of some strategic action. Tension in the model arises from a trade-off between reaping higher profits from moving later versus the possibility of preemption should other players move more quickly. Effectively, clock games are the oligopolistic analog of the competitive model of Abreu and Brunnermeier (2003). We showed that the unique symmetric equilibrium of the game has a remarkably simple structure—upon receiving a signal, each player optimally waits a fixed amount of time before moving. When moves of other players are observable, delay is still optimal so long as no player has moved. Once the first move is observed, however, all other players optimally move immediately.

We then tested the model using controlled laboratory experiments. We varied the observability of moves and the degree of information diffusion. We found broad support for the main implications of the Clock Games model and, more broadly, for the equilibrium effects of the trade-off between the waiting and preemption motives. Specifically, we observed considerable delay in the timing of moves by players after having observed the signal that the time to move was “ripe.” Indeed, this was the case even when moves were perfectly observable. We found that slower information diffusion led to longer delay while observable moves led to herding following the first move.

While many of the qualitative predictions of the model were borne out in the lab, several key discrepancies emerged between theory and observed behavior. First, unlike the theory model, we found that a longer calendar time before players received the signal that the time was “ripe” led to shorter delay and, in some cases, no delay at all. We found the force of the “calendar” effect differed depending on the speed of information diffusion. Finally, we found that the possibility of mistakes, especially when moves were observable, led to less herding behavior than was predicted by the theory. Taken together, these discrepancies suggest the need to modify the models to allow for the possibility of mistakes on the part of players, perhaps by analyzing the clock games framework using quantal response equilibrium as the solution concept. Given the complexity of the equilibrium characterization under full rationality, we felt it appropriate to leave this generalization for future research.

References

- ABREU, D., AND M. K. BRUNNERMEIER (2002): “Synchronization Risk and Delayed Arbitrage,” *Journal of Financial Economics*, 66(2-3), 341–360.
- (2003): “Bubbles and Crashes,” *Econometrica*, 71(1), 173–204.
- BAYE, M. R., D. KOVENOCK, AND C. G. DE VRIES (1993): “Rigging the Lobbying Process: An Application of the All-Pay Auction,” *American Economic Review*, 83, 289–294.
- BRUNNERMEIER, M. K., AND J. MORGAN (2006): “Clock Games: Theory and Experiments,” Working paper, Princeton University.
- BULOW, J., AND P. KLEMPERER (1999): “The Generalized War of Attrition,” *American Economic Review*, 89(1), 175–189.
- CABRALES, A., R. NAGEL, AND R. ARMENTER (2002): “Equilibrium Selection through Incomplete Information in Coordination Games: An Experimental Study,” *Universitat Pompeu Fabra Working Paper*, No. 601.
- CARLSSON, H., AND E. VAN DAMME (1993): “Global Games and Equilibrium Selection,” *Econometrica*, 61(5), 989–1018.
- CHAMLEY, C., AND D. GALE (1994): “Information Revelation and Strategic Delay in a Model of Investment,” *Econometrica*, 62, 1065–1085.
- FUDENBERG, D., AND J. TIROLE (1985): “A Theory of Exit in Duopoly,” *Econometrica*, 54(4), 943–960.
- (1991): *Game Theory*. MIT Press, Cambridge, MA.
- GUL, F., AND R. LUNDHOLM (1995): “Endogenous Timing and the Clustering of Agents’ Decisions,” *Journal of Political Economy*, 103, 1039–1066.
- HARRINGTON, J. E. (1989): “A Re-evaluation of Perfect Competition as the Solution to the Bertrand Price Game,” *Mathematical Social Sciences*, 17, 315–328.
- HARRIS, C., AND J. VICKERS (1985): “Perfect Equilibrium in a Model of a Race,” *Review of Economic Studies*, 52(2), 193–209.
- HEINEMANN, F., R. NAGEL, AND P. OCKENFELS (2004): “The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information,” *Econometrica*, 72(5), 1583–1599.
- HOPENHAYN, H., AND F. SQUINTANI (2006): “Preemption Games with Private Information,” *Working paper*.
- KUMMER, E. E. (1836): “Über die Hypergeometrische Reihe,” *Journal für reine und angewandte Mathematik*, 15, 39–83.

- LEININGER, W. (1989): “Escalation and Cooperation in Conflict Situations - The Dollar Auction Revisited,” *The Journal of Conflict Resolution*, 33, 231–254.
- LITTLEWOOD, J. E. (1953): *A Mathematician’s Miscellany*. Methuen & Co. Ltd, London.
- MCKELVEY, R. D., AND T. R. PALFREY (1992): “An Experimental Study of the Centipede Game,” *Econometrica*, 60(4), 803–836.
- MORRIS, S. (1995): “Co-operation and Timing,” *CARESS Working paper 95-05*.
- MORRIS, S., AND H. SHIN (1998): “Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks,” *American Economic Review*, 88(3), 587–596.
- OBSTFELD, M. (1996): “Models of Currency Crises with Self-Fulfilling Features,” *European Economic Review*, 40, 1037–1047.
- O’NEILL, B. (1986): “International Escalation and the Dollar Auction,” *The Journal of Conflict Resolution*, 30, 33–50.
- OSBORNE, M. J. (2003): *An Introduction to Game Theory*. Oxford University Press, Oxford.
- PARK, A., AND L. SMITH (2003): “Caller Number Five: Timing Games that Morph from One Form to Another,” *Mimeo, University of Michigan*.
- (2004): “From the Gold Rush to Rush Hour: Timing Models with Both Explosive and Timed Entry,” *Mimeo, University of Michigan*.
- REINGANUM, J. (1981): “Dynamic Games of Innovation,” *Journal of Economic Theory*, 25, 21–41.
- RIORDAN, M. (1992): “Regulation and Preemptive Technology Adoption,” *RAND Journal of Economics*, 23, 334–349.
- ROSENTHAL, R. (1981): “Games of Perfect Information, Predatory Pricing, and the Chain-Store Paradox,” *Journal of Economic Theory*, 25, 92–100.
- SHUBIK, M. (1971): “The Dollar Auction Game: A Paradox in Noncooperative Behavior and Escalation (in Gaming),” *Journal of Conflict Resolution*, 15(1), 109–111.
- SLATER, L. J. (1974): “Confluent Hypergeometric Functions,” in *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, ed. by M. Abramowitz, and I. A. Stegun, pp. 503–515. Dover, New York, NY, Chapter 13, 9th printing.
- WELFORD, A. T. (1980): *Reaction Times*. Academic Press, New York.
- ZHANG, J. (1997): “Strategic Delay and the Onset of Investment Cascades,” *Rand Journal of Economics*, 28, 188–205.

A Appendix

A.1 Properties of the Kummer Function

In this section, we detail some useful properties of the Kummer function, which we rely upon in what follows. One useful feature of this function is that it may be expressed as the infinite series

$$\Phi(a, b, x) \equiv \sum_{j=0}^{\infty} \frac{(a)_j x^j}{(b)_j j!} = 1 + \frac{a}{b}x + \frac{a(a+1)}{b(b+1)} \frac{x^2}{2!} + \dots,$$

where

$$(a)_k \equiv \begin{cases} 1 & \text{if } k = 0 \\ \prod_{n=0}^{k-1} (a+n) & \text{if } k > 0 \end{cases}.$$

From this representation, it may be readily seen that Φ is strictly increasing in its first and third arguments and strictly decreasing in its second argument. We will make extensive use of these monotonicity properties. In addition, the Kummer function has a number of other nice properties (see Kummer (1836)) which greatly simplify the analysis.

A.2 Proof of Proposition 1

A.2.1 Deriving hazard rate h_K

We will establish that the hazard rate of the K th lowest of $I - 1$ draws is independent of t_i . Formally $\frac{\Pr[t_K = t_i | t_i]}{\Pr[t_K \geq t_i | t_i]}$ is a constant.

First, fix t_0 . In that case, each player's type is uniformly distributed conditional on t_0 and hence the density of the K th order statistic is simply

$$f_K(t_K = t | t_i, t_0) = \frac{(I-1)!}{(K-1)!(I-K-1)!} \left(\frac{1}{\eta}\right)^{I-1} (t-t_0)^{K-1} (\eta - (t-t_0))^{I-1-K}$$

for $t \in [t_0, t_0 + \eta]$. Of particular interest is the case where $t = t_i$, in that case, we have

$$f_K(t_K = t_i | t_i, t_0) = \frac{(I-1)!}{(K-1)!(I-K-1)!} \left(\frac{1}{\eta}\right)^{I-1} (z_i)^{K-1} (\eta - z_i)^{I-1-K}, \quad (7)$$

where $z_i := t_i - t_0$. Of course, t_0 is not known to player i , but he knows t_i and hence that $t_0 \in [t_i - \eta, t_i]$. Given that t_0 is exponentially distributed, $\phi(t_0) = \lambda e^{-\lambda t_0}$,

$$\phi(t_0 | t_i) = \frac{\lambda e^{-\lambda t_0}}{1 - e^{-\lambda t_i} - (1 - e^{-\lambda(t_i - \eta)})},$$

which we may also write purely in terms of z_i to obtain

$$\phi(t_0 | t_i) = \frac{\lambda e^{\lambda z_i}}{e^{\lambda \eta} - 1}. \quad (8)$$

Consequently, we have that

$$f_K(t_K = t_i | t_i) = \frac{\lambda}{e^{\lambda\eta} - 1} \frac{(I-1)!}{(K-1)!(I-K-1)!} \left(\frac{1}{\eta}\right)^{I-1} \int_0^\eta e^{\lambda z_i} (z_i)^{K-1} (\eta - z_i)^{I-1-K} dz_i. \quad (9)$$

We also have to derive the denominator of the hazard rate, $\Pr[t_K \geq t_i | t_i]$. This is simply the probability that $K-1$ or fewer players received signals prior to t_i . Conditional on t_0 , these signals are uniformly distributed. Hence

$$f_K(t_K \geq t_i | t_i, t_0) = \sum_{n=0}^{K-1} \binom{I-1}{n} \left(\frac{1}{\eta}\right)^{I-1} (t_i - t_0)^n (\eta - (t_i - t_0))^{I-1-n}.$$

Again, it is convenient to express this in terms of differences z_i . Thus, we have

$$f_K(t_K \geq t_i | t_i, t_0) = \sum_{n=0}^{K-1} \binom{I-1}{n} \left(\frac{1}{\eta}\right)^{I-1} (z_i)^n (\eta - z_i)^{I-1-n}.$$

Integrating out t_0 , we obtain

$$f_K(t_K \geq t_i | t_i) = \frac{\lambda}{e^{\lambda\eta} - 1} \sum_{n=0}^{K-1} \binom{I-1}{n} \left(\frac{1}{\eta}\right)^{I-1} \int_0^\eta e^{\lambda z_i} (z_i)^n (\eta - z_i)^{I-1-n} dz_i. \quad (10)$$

Notice that both Equations 9 and 10 are independent of t_i . As a consequence, the hazard rate $h_K := \frac{\Pr[t_K = t_i | t_i]}{\Pr[t_K \geq t_i | t_i]}$ is also independent of t_i .

A.2.2 Deriving $E[e^{-g(t_i - t_0)} | t_i, t_K = t_i]$

From Bayes' rule, we know that

$$\phi(t_0 | t_i, t_K = t_i) = \frac{f_K(t_i | t_0) \phi(t_0 | t_i)}{f_K(t_i | t_i)}$$

From Equations (7) and (8), we know that $\phi(t_0 | t_i, t_K = t_i)$ may be rewritten entirely in terms of z_i . Now, notice that

$$E[e^{-g(t_i - t_0)} | t_i, t_K = t_i] = \int_{t_0 = t_i - \eta}^{t_i} e^{-g(t_i - t_0)} \phi(t_0 | t_i, t_K = t_i) dt_0.$$

may also be rewritten in terms of z_i as well. Hence,

$$E[e^{-g(t_i - t_0)} | t_i, t_K = t_i] = \frac{\int_0^\eta e^{-(g-\lambda)z_i} (z_i)^{K-1} (\eta - z_i)^{I-1-K} dz_i}{\int_0^\eta e^{\lambda z} (z_i)^{K-1} (\eta - z_i)^{I-1-K} dz_i},$$

which is a constant and hence, independent of t_i .

A.2.3 Simplifying τ

Rewriting Equation (2)

$$\tau = \frac{1}{g} \log \left(\frac{E [e^{-g(z_i)} | t_i, t_K = t_i]}{1 - \frac{g}{h}} \right).$$

and after cancellation, we can rewrite the above expression as:

$$\frac{E [e^{-g(z_i)} | t_i, t_K = t_i]}{1 - \frac{g}{h}} = \frac{\int_0^\eta e^{-(g-\lambda)z} (z)^{K-1} (\eta-z)^{I-1-K} dz}{\int_0^\eta e^{\lambda z} (z)^{K-1} (\eta-z)^{I-1-K} dz - g \int_0^\eta e^{\lambda z} \sum_{n=0}^{K-1} \frac{(K-1)!(I-1-K)!}{n!(I-1-n)!} (z)^n (\eta-z)^{I-1-n} dz}. \quad (11)$$

Now, let Num and Den denote, respectively, the numerator and denominator of the right-hand side of Equation (11). By series expansion of the exponential function it follows that

$$\begin{aligned} Num &= \frac{\eta^{I-1} \Gamma(I-K) \Gamma(K)}{\Gamma(I)} \Phi(K, I, \eta(\lambda - g)) \\ Den &= \frac{\eta^{I-1} \Gamma(I-K) \Gamma(K)}{\Gamma(I)} \left[\Phi(K, I, \eta\lambda) - \frac{\eta g}{I} \sum_{n=0}^{K-1} \Phi(1+n, 1+I, \eta\lambda) \right], \end{aligned}$$

where $\Gamma(\cdot)$ is the Gamma function satisfying $\Gamma(a) = (a-1)!$ for a positive integer a and $\Phi(\cdot)$ is the Kummer function defined above.

In the expression for Den ,

$$\begin{aligned} \sum_{n=0}^{K-1} \Phi(1+n, 1+I, \eta\lambda) &= \sum_{n=0}^{K-1} \sum_{j=0}^{\infty} \frac{(1+n)_j (\eta\lambda)^j}{(1+I)_j j!} \\ &= \sum_{j=0}^{\infty} \frac{1}{(1+I)_j} \frac{(\eta\lambda)^j}{j!} \sum_{n=0}^{K-1} (1+n)_j \\ &= \sum_{j=0}^{\infty} \frac{1}{(1+I)_j} \frac{(\eta\lambda)^j}{j!} \frac{(K)_{j+1}}{j+1}. \end{aligned}$$

The last equality follows from $\sum_{n=0}^{K-1} (1+n)_j = \frac{(K)_{j+1}}{j+1}$. Therefore,

$$\begin{aligned} \frac{\eta g}{I} \sum_{n=0}^{K-1} \Phi(1+n, 1+I, \eta\lambda) &= \frac{g}{\lambda} \sum_{j=0}^{\infty} \frac{(K)_{j+1} (\eta\lambda)^{j+1}}{(I)_{j+1} (j+1)!} \\ &= \frac{g}{\lambda} \sum_{j=1}^{\infty} \frac{(K)_j (\eta\lambda)^j}{(I)_j j!} \\ &= \frac{g}{\lambda} [\Phi(K, I, \eta\lambda) - 1]. \end{aligned}$$

Hence,

$$Den = \frac{\eta^{I-1} \Gamma(I-K) \Gamma(K)}{\Gamma(I)} \left[\left(1 - \frac{g}{\lambda}\right) \Phi(K, I, \eta\lambda) + \frac{g}{\lambda} \right].$$

Therefore,

$$\tau = \frac{1}{g} \log \frac{\lambda \Phi(K, I, \eta(\lambda - g))}{g - (g - \lambda) \Phi(K, I, \eta\lambda)}.$$

Next, we show that for $\eta \in [0, \bar{\eta}]$, $\tau > 0$. Since $E[e^{-g(z)} | t_i, t_K = t_i] > 0$, the following lemma is sufficient.

Lemma 2 $h_K > g$.

Proof. Since $\eta < \bar{\eta}$, $\lambda < g$, and Φ is increasing in its third argument, it then follows that

$$\begin{aligned} \frac{E[e^{-g(z)} | t_i, t_K = t_i]}{1 - \frac{g}{h_K}} &= \frac{\lambda \Phi(K, I, \eta(\lambda - g))}{g - (g - \lambda) \Phi(K, I, \eta\lambda)} \\ &> \frac{\lambda \Phi(K, I, \eta(\lambda - g))}{g - (g - \lambda) \Phi(K, I, \bar{\eta}\lambda)} \\ &= \frac{\lambda \Phi(K, I, \eta(\lambda - g))}{g - (g - \lambda) \left(\frac{I g}{I g - (I - K + 1) \lambda} \right)} \\ &> \Phi(K, I, \eta(\lambda - g)) > 0. \blacksquare \end{aligned}$$

A.3 Proof of Proposition 2

Recall that $\tau = \frac{1}{g} \log \frac{\lambda \Phi(K, I, \eta(\lambda - g))}{g - (g - \lambda) \Phi(K, I, \eta\lambda)}$. We first show that when $\eta = 0$, $\tau = 0$. From the series expansion of $\Phi(\cdot)$ it can easily be seen that $\Phi(K, I, 0) = 1$, hence when $\eta = 0$, $\tau = 0$.

For $\eta \in (0, \bar{\eta})$,

$$\frac{\lambda \Phi(K, I, \eta(\lambda - g))}{g - (g - \lambda) \Phi(K, I, \eta\lambda)} = \frac{1 + \int_0^\eta \frac{\partial}{\partial w} \Phi(K, I, w(\lambda - g)) \Big|_{w=z} dz}{1 - \frac{g-\lambda}{\lambda} \int_0^\eta \frac{\partial}{\partial w} \Phi(K, I, w\lambda) \Big|_{w=z} dz}.$$

Using the fact that $\frac{\partial}{\partial x} \Phi(a, b, x) = \frac{a}{b} \Phi(a+1, b+1, x)$, it follows that the right-hand side may be rewritten as

$$\begin{aligned} &= \frac{1 - \frac{(g-\lambda)K}{I} \int_0^\eta \Phi(K+1, I+1, z(\lambda - g)) dz}{1 - \frac{(g-\lambda)K}{I} \int_0^\eta \Phi(K+1, I+1, z\lambda) dz} \\ &= 1 + \frac{\frac{(g-\lambda)K}{I} \int_0^\eta \Phi(K+1, I+1, z\lambda) - \Phi(K+1, I+1, z(\lambda - g)) dz}{1 - \frac{(g-\lambda)K}{I} \int_0^\eta \Phi(K+1, I+1, z\lambda) dz}. \end{aligned}$$

That this expression is increasing in η follows from the fact that $\Phi(\cdot)$ is increasing in its third argument. ■

A.4 Proof of Lemma 1

Let $t' > t_0$ be the time at which the first exit occurred. Since the time of the first exit is publicly observed, it becomes common knowledge that $e^{gt'} > e^{gt_0}$. The continuation game is strategically equivalent to pure Bertrand competition with a stochastic outside option whose highest possible realization lies strictly below the minimum possible bid. A straightforward extension of Harrington (1989) shows that, the unique equilibrium in this game is for all players to bid the minimum amount, $e^{gt'}$, i.e. all types exit immediately.

A.5 Proof of Proposition 3

It is straightforward to obtain expressions for h_1 and $E[e^{-g(z_i)}|t_i, t_1 = t_i]$. Simply use the analogous expressions given in the proof of Proposition 1 and substitute $K = 1$. This yields

$$\tau_1 = \frac{1}{g} \log \left(\frac{\int_0^\eta e^{-(g-\lambda)z} (\eta-z)^{I-2} dz}{\int_0^\eta e^{\lambda z} (\eta-z)^{I-2} dz - \frac{Ig}{I-K+1} \int_0^\eta e^{\lambda z} \frac{(I-2)!}{(I-1)!} (\eta-z)^{I-1} dz} \right).$$

Using steps analogous to the simplification of τ , we have

$$\tau_1 = \frac{1}{g} \log \left(\frac{\lambda \Phi(1, I, \eta(-g+\lambda))}{\frac{Ig}{I-K+1} - \left(\frac{Ig}{I-K+1} - \lambda\right) \Phi(1, I, \eta\lambda)} \right).$$

It remains only to show that $\tau_1 > 0$. Since $E[e^{-g(z)}|t_i, t_1 = t_i] > 0$, the required inequality follows from the following Lemma.

Lemma 3 $h_1 > \frac{Ig}{I-K+1}$.

Proof. Following the identical steps in the proof of Lemma 2, we obtain

$$\frac{E[e^{-g(z)}|t_i, t_1 = t_i]}{1 - \frac{g_1}{h_1}} = \frac{\lambda \Phi(1, I, \eta(-g+\lambda))}{\left(\lambda - \frac{Ig}{I-K+1}\right) \Phi(1, I, \eta\lambda) + \frac{Ig}{I-K+1}} \geq 0, \quad (12)$$

which is satisfied since $\eta \leq \bar{\eta}$. ■

Thus, we have shown that τ_1 strategies comprise a symmetric equilibrium. The fact that τ_1 is the unique symmetric equilibrium follows using steps identical to those in Proposition 1. ■

B Appendix

Note: Terms in {.} are included only in the COMPRESSED treatments while terms in [.] are included only in OBSERVABLE treatments.

Instructions

Thank you for participating in this experiment on the economics of investment decision making. If you follow the instructions carefully and make good decisions you can earn a considerable amount of money. At the end of the experiment you will be paid in cash and in private. The experiment will take about an hour and a half.

There are 12 people participating in this session. They have all been recruited in the same way that you have and are reading the same instructions that you are for the first time. Please refrain from talking to the other participants during the experiment.

You are about to play the same selling game 45 times in succession. Each round represents one trading period and each person in the room is a seller in the game. There are two games running simultaneously, so you will not know which game you are in and which players in the room you are playing with. The sellers are randomly matched at the beginning of each period and thus the composition of sellers for each game changes from trading period to trading period.

How to play the game:

Initially, everyone needs to login from the login page. The game begins after five other players log in (see Figure 1). At the start of each trading period, the price of everyone's asset begins at 1 ECU (experimental currency units) and increases exponentially. The true value of the asset is predetermined when the game starts. At one point, you will be notified that the current price of the asset exceeds its true value. At that time, the price of the asset will change from red to green. In addition, you will be given the minimum and maximum values of the asset. The true value of the asset lies between those two values (see Figure 2).

[In addition, whenever any other seller in your game chooses to sell, you will be notified of this fact by a message at the bottom of your screen indicating the price at which the sale occurred.]

Please note that the speed to which the price rises will vary slightly. This is due to random network traffic.

How to Sell: In each period, you make your decision of selling your asset by moving the mouse pointer into blue “Sell” box on your screen. DO NOT CLICK THE MOUSE BUTTON. Once you have decided to sell, you have no more decisions to make in the trading period. Your decision is final.

There are two ways a trading period can end:

- (1) Once the third seller moves his/her mouse into the blue window, the period ends;
- (2) If fewer than three sellers sell, then the trading period ends once a pre-determined number of seconds has elapsed after the asset reaches its true value.

If you sell, you earn an amount in ECUs equals to the price you sell at. If you don't sell, you will earn the true value of the asset. Please keep in mind that your goal is to maximize your earnings.

At the end of each trading period, the true value of the asset, your earnings for this period, your cumulative earnings, and the earnings of other players will all be displayed on your screen (see Figure 3).

Click the “Play Again” link on the screen to play the next round. The new round will begin when every player has clicked the link.

At the conclusion of the experiment, your total earnings in ECUs will be converted to cash at the rate of 50 ECU = \$1.

At some point during the trading period, you will receive a message indicating that the current price of the asset has exceeded its “true value”. When you receive this message, the computer will also inform you of the minimum and maximum “true values” for the asset.

Are there any questions?

Details – This section contains technical details about the game

At the start of each selling game, the price of the asset begins at 1 ECU (experimental currency units). The price of the asset increases by 2% in each trading period (which lasts about 1 second). You will see the price increasing on your computer screen---the current price is the same for all sellers.

In addition, the computer secretly determines the *true value* of the asset. The shadow value increases with the price of the asset until the computer determines that it has stopped growing. In each trading period, there is a 1% chance that the true value will stop growing. The true value of the asset is the same for you and the other five sellers with whom you are participating.

Delayed Information

Once the computer determines that the true value has stopped growing, you will be alerted of this fact, **but** with a random delay. The computer will choose a random delay from zero, one, two, up to ninety {fifty} trading periods (equally likely) for each seller. So on average, the delay will be about 45 {25} seconds from the time the shadow value stopped growing until the time you become informed of this fact. After this delay, you will see a message on your computer screen indicating that the true value has stopped growing as well as displaying the highest and lowest amounts the shadow value could be. The same is true for the other five players, but the computer determines their delay separately your delay. That is, most likely they hear the news at different points in time.

End of the Game

Each selling game ends after three people have sold the asset or 200 periods after the true value of the asset has stopped growing, whichever comes first.

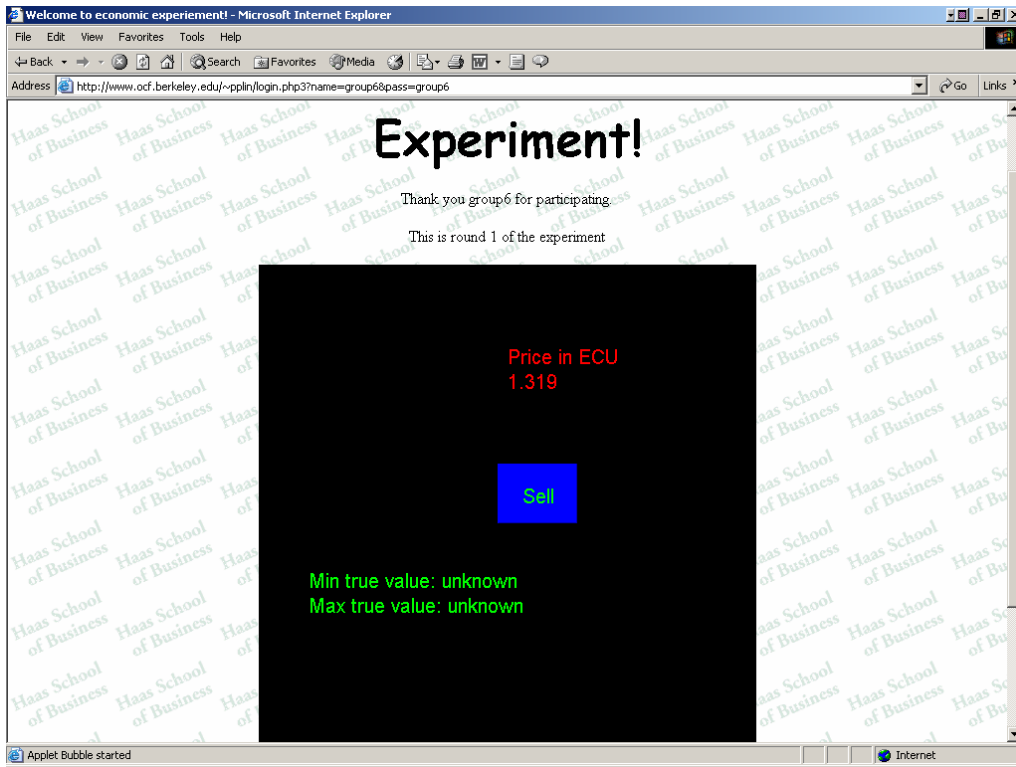


Figure 1

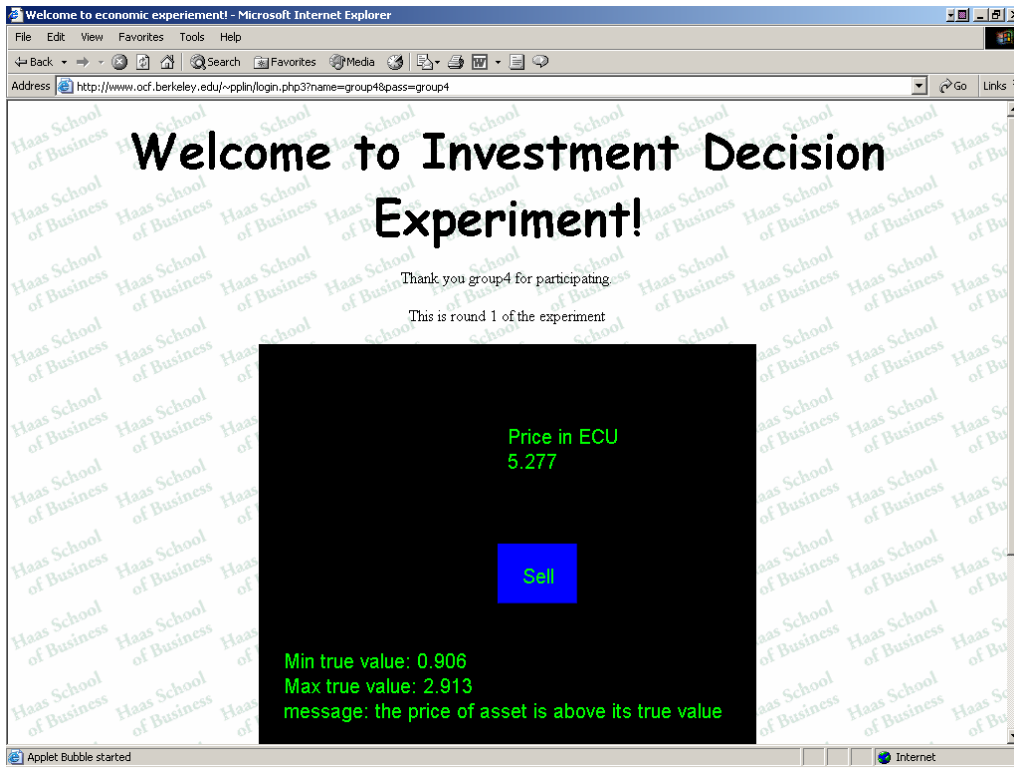


Figure 2




Figure 3