

UC Berkeley

Recent Work

Title

Clock Games: Theory and Experiments

Permalink

<https://escholarship.org/uc/item/9c11m09n>

Authors

Brunnermeier, Markus K
Morgan, John

Publication Date

2006-10-03

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: October 3, 2006

Abstract

Timing is crucial in situations ranging from product introductions, to currency attacks, to starting a revolution. These settings share the feature that payoffs depend critically on the timing of moves of a few other key players—and these are uncertain. To capture this, we introduce the notion of clock games and experimentally test them. Each player’s clock starts on 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) equilibrium is always characterized by strategic delay—regardless of whether moves are observable or not; (ii) delay decreases as clocks become more synchronized and increases as information becomes more concentrated; (iii) When moves are observable, players “herd” immediately after any player makes a move. We then show, in a series of experiments, that key predictions of the model are consistent with observed behavior.

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.

[†]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

In many business situations, timing is critical to succeed in the marketplace. For instance, a key decision firms face is when to introduce a new product. Launching too soon, either before the technology is sufficiently developed or before the consumer is ready to accept the product, can lead to difficulties. On the other hand, waiting too long may be costly as well, owing to rivals' filling the product space in the interim. A concrete example is the market for MP3 players. One early entrant to this market, the Diamond Rio player, was the market share leader early on but did not achieve widespread consumer acceptance. Apple's iPod, a later entrant, soon displaced the Diamond Rio player and other earlier entrants. Apple continues to dominate the MP3 player space currently thanks to the maturation of hard drive miniaturization technology and an agreement for the legal distribution of digital music content with major providers. A host of iPod "wannabes" followed—including offerings from Creative Labs, Dell, SanDisk, and many others. Despite being technologically similar to the iPod, these later rivals have had little success in gaining market share at Apple's expense. Entry now is very unlikely to be profitable and indeed some of the existing players, like Rio, are now starting to exit this space.

Outside of a business context, many other decisions exhibit similar trade-offs. For instance, consider the situation of a political revolution. Early revolutionary leaders are unlikely to be successful if the existing regime still possesses sufficient capabilities to "quiet" dissidents. As popular momentum for the revolution increases, revolutionary leaders are more likely to be successful and to gain positions of political power in the successor regime. Entering too late, after the power vacuum is filled, is also unlikely to lead to success. A third class of examples are currency attacks: An "attacker" who moves too early will be countered by the central bank with high interest rates. This leads to less profitable outcomes and even losses if the attack cannot be sustained. Moving too late is equally problematic as other investors may already have "attacked" the currency and erased the mispricing.

A situation perhaps more familiar to the average reader where timing is crucial concerns investing in the stock market. Those who missed the incredible run-up of tech sector stocks in the US in the late 1990s had good cause to regret their bad timing. Likewise, those who stayed in the market too long and suffered the large downturn of the early 2000s also had good cause to regret their inaction.

All of these situations have in common the presence of both a waiting motive—with patience the move becomes more valuable—and a preemption motive—move too late and a rival will reap the rewards. The timing of the rivals' moves represents the key strategic uncertainty faced by players. Yet, the situations described above differ from one another in other key respects. Most notably, the situations above differ in the observability of "rival" actions. In the case of new product introductions, the moves by Apple and its rivals were, arguably, observable (although perhaps with a lag owing

to R&D lead time) and thus provided an opportunity to react. Moreover, many of these situations share the feature that there are only a few key players, each of whom is “large” that affect outcomes. For instance, in situations of political revolutions, there are typically a few key players, such as Napoleon, Lenin, or Khomeini, with the charisma and resources needed to effect regime change. How is the timing of strategic moves affected by observability and the number and informational size of players?

To examine these questions we analyze a class of games we refer to as “clock games.” Our analysis consists of both a theoretical examination as well as testing through controlled laboratory experiments. 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 (i.e., the time is ripe for a product introduction, etc.). 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 number and size of players, and the speed of information diffusion.

The nearest theoretical antecedent to our paper is Abreu and Brunnermeier (2003), hereafter referred to as AB. Like AB, we also study equilibrium delay when waiting and preemption motives are present. Unlike AB, our framework allows for variation in (i) the number of players; (ii) the transparency of moves of players; and (iii) variation in the informational “size” of players participating in the game and nests the AB model as a special case. The generalizations of the AB model contained in the clock games framework are important for two reasons. First, the AB model, in its original form, is inherently untestable in the field owing to its informational assumptions or the laboratory owing to its assumption that there are a continuum of players. While tests using field data are difficult to do convincingly for all models of this type, by allowing for a discrete number of competing agents, the clock games framework restores the testability of these models in the laboratory. Second, some of the predictions of the AB model are sensitive to its particular specification. Indeed, when moves are observable, the AB model has the perverse implication that the preemption-waiting tradeoff vanishes—equilibria do not entail delay. Our framework shows that this is an artifact. The preemption-waiting tradeoff—and resulting equilibrium delay—remains even with observable moves.

To summarize, the main contributions of this paper are (i) to generalize the AB model to account for the possibility that strategic moves are observable or players are “large”; and (ii) to test the empirical validity of the predictions of the this class of models in a controlled laboratory setting.

Our main results concerning the theory of clock games are:

1. Equilibrium delay is robust—it occurs regardless of observability of moves, number, or informational size of players.
2. When moves are observable, equilibrium delay occurs *only* when individuals are

few or informationally large. After the first player moves, equilibrium entails “herding”—all other players move immediately after observing that move. With a continuum of informationally small players (as in the AB model), there is still herding but *no equilibrium delay* when moves are observable.

3. The slower is information diffusion, the longer is the equilibrium delay.
4. The fewer or informationally larger are the players, the longer is equilibrium delay.

To test the theory, we then 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.

Our main results concerning experiments of clock games 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 is information diffusion, the longer is the observed delay.

While the three results reported above are highly consistent with the theoretical predictions of clock games models, we observe a number of systematic differences. First, there is considerable heterogeneity in behavior. While the theory predicts that individuals will move a fixed number of periods after receiving the signal, the experiments reveal that some subjects systematically move earlier than others. Second, in some instances this takes an extreme form: some subjects move prior to receiving the signal at all—a move not predicted by the theory. We find that this behavior is correlated with the timing of the payoff relevant state variable—the later in time this variable appears, the greater the propensity for subjects to exit early. Another complication which might explain discrepancies between the theory and observed behavior is that mistakes are likely to occur in actual behavior. Thus, the decision making process of a subject in the experiment is made more difficult by the fact that he or she must account for mistakes of others, which greatly complicates the inference problem. We show that many of the discrepancies between theory and observed behavior may be rationalized if subjects follow heuristic strategies but misperceive (in a natural way) the statistical process governing the timing of the payoff-relevant state variable.

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. In Section 2 we present the model of clock games, characterize equilibrium play, and identify key testable implications of the model. Section 3 outlines the procedures used to test the theory in controlled laboratory experiments. In Section 4 we present the results of

the experiments, comparing and contrasting these results with the predictions of the model. Section 5 presents a discussion of the results of the experiment and suggests directions in which the theory might be amended based on the behavioral observations of the experiment. Section 6 then extends the basic clock games model to situations where there are a continuum of players; thus, nesting the AB model in the clock games framework. Finally, Section 7 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.

Pre-emption 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 backward 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 backward induction plays in the two games. Since there is no *commonly known* point from which one could start the backwards induction argument, the backward induction rationale does not appear in clock games with unobservable moves whereas it is central in centipede games. Clock games differ in another respect as well: In clock games, multiple players can move simultaneously, while in the centipede game the two players alternate.

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 to 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.

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.

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 information, our concerns center on the role of private information and, in particular, 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 a generalization of the class of models in Abreu and Brunnermeier (AB 2002, 2003). AB (2002, 2003). These papers study persistence of mispricing in financial markets with a continuum of informationally small, anonymous traders.

2 Theory

2.1 Model

We study the following situation: A finite number, I , of players are participating in a game analogous to the situations described in the introduction. At the start of the game, each of the players is currently “in” the game and the only decision is when to exit. Once a player exits, he or she cannot subsequently return; thus, each player’s strategy amounts to a simple stopping time problem.³

The game can end in one of two ways: First, the game ends at the point in time

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

³Consistent with standard models in the stopping time literature, we treat time as continuous. This allows us to obtain tractable closed-form solutions. Technically, this assumption is at odds with the experimental implementation (where time is necessarily discrete). To ensure the robustness of our theoretical findings to the experimental setting, we computed numerically equilibrium stopping time strategies for the discrete version of the model and verified that they converged to those for the continuous time case in the limit.

when a critical number, $K < I$, of the players have exited.⁴ Second, the game ends at time $t_0 + \bar{\tau}$, where $\bar{\tau}$ is commonly known, if fewer than K players exited by this point. A player's payoff is determined by whether or not he exits before the game ends. If he exits before the game ends, at time t (say), then his "exit" payoff is e^{gt} . If, however, the game ends before the player exits, then he receives an "end-of-game" payoff, e^{gt_0} . The random variable t_0 corresponds to the time in which the exit payoff starts to exceed the end-of-game payoff. We assume that t_0 is exponentially distributed with p.d.f. $f(t_0) = \lambda e^{-\lambda t_0}$, where the constant λ is the arrival rate that t_0 occurs in the next instant conditional on the event that it did not happen so far. Finally, if one player exits exactly when the game ends, he still receives the exit payoff. If however more players exit exactly at this point, we employ the following tie-breaking rule: Suppose that up to this point $L < K$ players have exited. Then each of the remaining players has an equal chance of obtaining one of the remaining $K - L$ available "slots" and obtaining the exit payoff.




Figure 1: 'Exit payoff' versus 'end of game payoff'

Interest in the model arises from the fact that a player can suffer drop in payoff by waiting too long to exit. That is, for any time beyond t_0 , a player's end-of-game payoff is lower than his payoff if he exited at that moment since by waiting another instant and exiting, the player enjoys growth rate g provided the game does not end in the interim. Figure 1 illustrates both types of payoffs for the case where $t_0 = 130$

⁴As we explain in detail below, the game also ends if fewer than K players exit and the game has been ongoing for a sufficiently long time.

and $g = 2\%$. The payoffs from exiting (solid line) lie below the end-of-game payoffs (dotted line) for $t < t_0$ and above for $t > t_0$.

At time $t_i \geq t_0$, player i receives a signal indicating that the exit payoff exceeds the end-of-game payoff. It is helpful to think of time t_i as the point at which player i 's clock starts. The timing of i 's signal is uniformly distributed within the “window of awareness” $[t_0, t_0 + \eta]$, where η is the length of the window. That is, each player does not exactly know when others received their signals. For instance, player i , who receives a signal at time t_i , only knows at time $t_i + \eta$ that all other players received signals. In Figure 1 the shaded rectangle illustrates the window of awareness for the case where $\eta = 50$.

For the model to be interesting, the following assumptions are sufficient: (i) $0 < \lambda < g$, (ii) $\bar{\tau}$ large and (iii) η not too large. Assumption (i) guarantees that there is sufficient upside to waiting, and so strategic delay becomes a possibility. Assumption (ii) ensures that the prospect of the game ending for exogenous reasons is not a strategic consideration. Finally, assumption (iii) 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. Assumption (iii) may be stated more precisely as follows: Let $\bar{\eta}$ solve $F(K, I, \bar{\eta}\lambda) = \frac{Ig}{Ig - (I - K + 1)\lambda}$, where the function $F(a, b, x)$ is a Kummer hypergeometric function (see e.g. Slater (1974)).⁵ From the monotonicity properties of $F(\cdot)$, such a solution always exists and is unique. Assumption (iii) requires that $0 < \eta < \bar{\eta}$.

While the model seeks to capture the central features shared by the situations described in the introduction where both waiting and preemptive motives are present and where strategic uncertainty about the timing of a rival's move is critical, it abstracts away from detailed features unique to each situation. For instance, we do not model a change in the profit path of new product introductions as a function of the timing of past introductions other than to allow for a discontinuity in prices once the K th entry occurs. Moreover, continuous effort in the form of the quality of the new product offering, the financial size of the currency attack, nor the size of the revolutionary organization are modeled. Instead, we view our model as a basic model useful for developing implications that are testable in the lab and as providing a kind of baseline model which might be enhanced for detailed study of any of the situations described in the introduction.

Next, we characterize symmetric perfect Bayesian equilibria for two cases of the

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

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

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

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

Since all of the players in the game are *ex ante* identical, we restrict attention to symmetric equilibria. In Proposition 1 we show that there is a unique symmetric equilibrium in our game. In this equilibrium, each player waits exactly τ periods after receiving his or her signal before exiting the game and exits immediately thereafter. We present a heuristic proof to illustrate the construction of this equilibrium below.⁶ In the Appendix, we formally establish both existence and uniqueness.

First, fix the strategies of all other players as described above and consider the problem faced by player i at time $t_i + \tau$. Player i faces an endogenous hazard rate, h , associated with the chance that the game will end in the next instant. For player i to decide to exit the game at time $t_i + \tau$ rather than to stay in, it must be the case that player i 's expected profit from exiting at $t_i + \tau$ is more than the expected profit from exiting at $t_i + \tau + \Delta$. For small Δ , we can focus on the linear approximation of i 's payoffs ignoring the tie-breaking rule. Thus, the change in expected profit from delaying an additional Δ periods is

$$(1 - h\Delta) g e^{g(t_i + \tau)} \Delta - h\Delta E [e^{g(t_i + \tau)} - e^{gt_0} | D_\Delta, t_i].$$

With a probability of approximately $(1 - h\Delta)$ the payoff increases at a rate of g . Note that the expectations are taken conditional on the fact that the game will end in the next Δ interval (D_Δ) and on the time when i received the signal (t_i). With probability of (approximately) $h\Delta$ the payoff drops (i.e., the game ends) within $t_i + \tau$ and $t_i + \tau + \Delta$ and exiting at $t_i + \tau$ leads to the higher exit payoff $e^{g(t_i + \tau)}$ rather than the end-of-game payoff e^{gt_0} . Letting Δ go to zero, the second order Δ^2 -terms vanish and, the optimal stopping time equates the marginal (log) benefits of delaying exiting with marginal (log) costs. That is,

$$hE [1 - e^{-g(t_i + \tau - t_0)} | D_0, t_i] = g. \quad (1)$$

Note that $[1 - e^{-g(t_i + \tau - t_0)}]$ is the drop in payoff as a fraction of the current payoff $e^{g(t_i + \tau)}$. Solving for τ in equation (1) yields the optimal stopping time for player i

$$\tau = \frac{1}{g} \left[\ln \frac{h}{h - g} + \ln E [e^{-g(t_i - t_0)} | D_0, t_i] \right]. \quad (2)$$

Of course, equation (2) depends on the hazard rate, h , as well as on the conditional expectation $E [e^{-g(t_i - t_0)} | D_0, t_i]$. Both expressions are determined in equilibrium.

⁶The proof is heuristic because we restrict attention to local deviations.

Let us consider these terms in more detail. At time $t_i + \tau$ player i 's likelihood that the payoff will drop in the next instant is as likely as the event that the K th of the other $I - 1$ players receives his signal exactly at t_i (provided that $t_i + \tau < t_0 + \bar{\tau}$). For this event to occur, $K - 1$ of the other players receive their signals prior to t_i , while one player, the K th to receive the signal, gets this information exactly at t_i , and the remaining $I - 1 - K$ players receive signals after t_i . To derive the hazard rate, we specify in the Appendix player i 's likelihood that the game will end in the next instant for a given t_0 . Second, we take expectations over different t_0 realizations and finally, we condition on the event that the game is ongoing. This yields an equilibrium hazard rate of:

$$h = \frac{\frac{(I-1)!}{(K-1)!(I-1-K)!} \int_0^\eta e^{\lambda z} z^{K-1} (\eta - z)^{I-1-K} dz}{\int_0^\eta e^{\lambda z} \sum_{n=0}^{K-1} \frac{(I-1)!}{n!(I-1-n)!} z^n (\eta - z)^{I-1-n} dz}. \quad (3)$$

Note that the hazard rates are constant in equilibrium due to the exponential prior distribution of t_0 .

After deriving the distribution of t_0 conditional on a drop in payoff the instant after player i exits, we show in the Appendix that

$$E [e^{-g(t_i - t_0)} | D_0, t_i] = \frac{\int_0^\eta e^{-(g-\lambda)z} z^{K-1} (\eta - z)^{I-1-K} dz}{\int_0^\eta e^{\lambda z} z^{K-1} (\eta - z)^{I-1-K} dz}. \quad (4)$$

Substituting equations (3) and (4) into equation (2) yields

$$\tau = \frac{1}{g} \ln \left(\frac{\int_0^\eta e^{-(g-\lambda)z} z^{K-1} (\eta - z)^{I-1-K} dz}{\int_0^\eta e^{\lambda z} z^{K-1} (\eta - z)^{I-1-K} dz - g \int_0^\eta e^{\lambda z} \sum_{n=0}^{K-1} \frac{(K-1)!(I-1-K)!}{n!(I-1-n)!} z^n (\eta - z)^{I-1-n} dz} \right).$$

This unwieldy integral expression is greatly simplified using Kummer functions, $F(a, b, x)$.

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

$$\tau = \frac{1}{g} \ln \left(\frac{\lambda F(K, I, \eta(\lambda - g))}{g - (g - \lambda) F(K, I, \eta\lambda)} \right). \quad (5)$$

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




Figure 2: Equilibrium delay, τ , for different η

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 $\lambda = 1\%$.

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 5, 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 the I , the number of players.

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.

The following observation is crucial: Suppose that there is an equilibrium where, prior to anyone exiting, players exit only *after* having received the signal, then on observing the first exit, the maximum time that can elapse before the game ends becomes common knowledge for the remaining players.⁷ Thus, if the first exit occurs at time t_1 , then it is common knowledge that the game will end no later than at time $t_1 + \bar{\tau}$. Intuitively, the presence of a commonly known finite ending time to the game now allows one to apply backward induction in an analogous fashion to a number of other timing games.⁸ A straightforward implication of this observation is the following:⁹

Proposition 3 *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.*

A key testable implication of Proposition 3 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, Propositions 1 and 3 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. To derive the equilibrium delay, τ_1 , let us again consider player i at time $t_i + \tau_1$. If he delays exiting by an additional Δ interval, he gains approximately $ge^{g(t_i + \tau_1)}\Delta$ if the game does not end. This event occurs with probability approximately $1 - \Delta h_1$, where h_1

⁷We make the usual assumption that all players’ conjectures about equilibrium strategies are commonly known.

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).

⁹While the argument above constitutes a proof for the discrete time case, our continuous time modeling raises technical issues in applying backward induction. We offer a formal proof of Proposition 3 in the Appendix.

is the endogenous hazard rate that some other player exits in the next instant. From Proposition 3, we know that, in this event, all remaining players will exit immediately in the subsequent instant. Following our tie-breaking rule, each exiting player has equal chance $\frac{K-1}{I-1}$ of receiving the exit payoff $(1 + g\Delta) e^{g(t_i+\tau_1)}$. Otherwise, a player receives only the end-of-game payoff, e^{gt_0} . This yields the first-order condition:

$$(1 - \Delta h_1) g e^{g(t_i+\tau_1)} \Delta + \Delta h_1 \left\{ \frac{K-1}{I-1} g e^{g(t_i+\tau_1)} \Delta - \frac{I-K}{I-1} E [e^{g(t_i+\tau_1)} - e^{gt_0} | D_\Delta, t_i] \right\} = 0.$$

As Δ goes to zero, the second order Δ^2 -terms vanish, and the first-order condition simplifies to

$$\tau_1 = \frac{1}{g} \left[\ln \frac{h_1}{h_1 - g \frac{I-1}{I-K}} + \ln E [e^{-g(t_i-t_0)} | D_0, t_i] \right]. \quad (6)$$

Comparing equation (6) 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 (2). 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 that given in equation (3) if one sets $K = 1$. Finally, note that the term $E [e^{-g(t_i-t_0)} | D_0, t_i]$ 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 4 *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*

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

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

Proposition 4 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.

How do the equilibrium delay times compare in the observable versus unobservable cases? In general, the effect 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. We cannot rule out the possibility that framing the clock game as a product introduction or even a purely abstract situation will lead to different behavior. Nonetheless, our view is that this frame offers a reasonable starting point for examining the behavioral relevance of some of the main predictions arising from our clock games model.

In our design each session consisted of 45 “rounds” or iterations of the game, all under the same treatment.¹⁰ 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 each. The job of a trader was to decide at what price to sell the asset they were holding. In making this decision, 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. Specifically the price of the asset was updated twice per second (in real time). The price of the asset began

¹⁰Owing to networking problems, session 3 lasted only 35 rounds.

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. The computer also determined in which period the “true value” of the asset had stopped growing. There was a 1% chance of this event for each period ($\lambda = 1\%$). In addition, 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 given market sold their unit of the asset, the round ended and each subject learned his or her earnings for the round and cumulative earnings for the experiment to date.¹¹ Each subject also learned the prices at which all of the assets sold in their market. The earnings of a subject in a given 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). In terms of the theory model, all experimental sessions used the parameter values $I = 6$, $K = 3$, $g = 2\%$, and $\lambda = 1\%$. We also set $\bar{\tau} = 200$.

At the end of the session, subjects were paid at the exchange rate of 50 ECUs to \$1, with fractions rounded up to the nearest quarter. Earnings averaged \$15.16 and each session lasted from 50 to 80 minutes.

Treatments Of central interest is how changes in both observability and the window of awareness impact the timing of exit decisions. That is, the experimental treatments are designed to test the main implications of Propositions 2 through 4. To examine these implications, we ran sessions under three different treatments: In the Baseline treatment, we set the window of awareness, $\eta = 90$. That is, each subject learned that the price of the asset exceeded the true value with a delay time that was uniformly (and independently) distributed from 1 to 90 periods following the event that the true value of the asset stopped growing. In the Compressed treatment, we reduced the window of awareness, η , from 90 to 50. Thus, comparing behavior in Baseline versus Compressed treatments allows us to test Proposition 2.

Finally, in the Observable treatment the window of awareness was the same as in Baseline; however, subjects received messages indicating each time a trader sold an asset in the market. That is, trading information was observable. Thus, comparing behavior in Baseline versus observable treatments allows us to test herding behavior (Proposition 3) as well as comparing the length of strategic delay (Proposition 4).

We ran six sessions each under the Baseline and Compressed treatments and four sessions under the Observable treatment giving 16 sessions overall. A total of 192

¹¹In 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.

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. That is, we wanted to minimize the possibility that subjects might use various auditory “cues” in detecting trading behavior by other subjects in the Baseline and Compressed treatments. Specifically, we were concerned that having the subjects click their mouse on the sell button in order to sell would enable other traders to detect selling by listening for mouse clicks. To remedy this problem, our experimental design had subjects sell by *hovering* their mouse over the sell button.

This was very effective in minimizing information leakage. However it did occasionally lead to subjects making what appear to us to be selling “mistakes.” Many subjects evolved the strategy of placing their mouse pointer close to the sell box so that they could quickly sell, but occasionally, their mouse pointer would inadvertently stray into the sell box resulting in an unintended sale. Many of these mistakes are fairly obvious in the data in that the first sale would sometimes take place after extremely few periods has occurred after the start of the round. In the case of the Baseline and Observable treatments, we “cleaned” the data by eliminating observations where sales occur 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 somewhat complicated. Thus, we expected that subjects would require several rounds of “learning by doing” before converging to a strategy as to how to play the game. As a consequence, our design stressed repetition in a stationary environment (45 iterations of the same treatment). We also tried to speed the learning process by giving subjects extensive feedback about the profitability of their decisions as well as a comparison group consisting of the profits of other traders in the same market. Finally, when a subject sold the asset *below* its true value, the subject received a message that this was the case in addition to his or her usual report about trading profits at the end of a round.

We observed considerable variability in subject choices in the early rounds of the game, suggesting that subjects were still learning the game. Behavior displayed much less variability in the last 25 rounds of each session. Since we are primarily interested in the performance of the model in equilibrium and not in learning to play the equilibrium strategy, we confine attention in the results section below to these periods.¹²

One worry we had about running 45 iterations was that the game would become, in effect, a repeated game for subjects. To counter this possibility, we randomly and

¹²This 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.

anonymously rematched subjects into different groups after each round of the game. Further, we prohibited communication among subjects. Thus, 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 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. Thus, we are reasonably confident that subjects were, in fact, treating the game as a one-shot game as described in the theory.

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.¹³

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 no *Delay* is assigned. If i sold at or before the time he received the signal, *Delay* equals zero.¹⁴
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 4—namely that equilibrium behavior will lead traders to engage

¹³The formal empirical analysis makes no use of the answers given in the questionnaire.

¹⁴We 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.

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 Proposition 3—that observable trading information leads to “herding” on the part of sellers following the first sale.¹⁵ Table 1 presents the predictions of the theory model for each of these performance measures. The parameters chosen for each of the treatments were designed to generate large differences in the performance measures. In particular, 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. Finally, the expectation of the Gap measure illustrates a distinct difference between the Baseline and Compressed treatments and the Observable treatment, which is predicted to have negligible gap length.

Table 1: Theory Predictions

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

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

4.1 Overview

Table 2 presents descriptive statistics from the experimental data for these same performance measures, treating each session as an independent observation. Beginning with the Duration measure, notice that the ordering implied by the theory is mainly reflected in the data: the longest average Duration occurs in the Baseline treatment, whereas the Observable and Compressed treatments exhibit shorter, but comparable Durations.

¹⁵One 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).

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

Turning to the Delay measure, notice that for all sellers, the ordering of Delay also mirrors the comparative static predictions of the theory. Delay is longer in Baseline sessions, followed by Observable sessions, followed by Compressed. However, an ancillary prediction of the theory—that Delay should be independent of the identity of the seller (i.e., first, second or third seller) is not borne out in Table 2. Indeed, there seems to be a consistent pattern that the first seller’s delay is shorter than the second, whose delay in turn is shorter than the third seller’s delay. We turn to possible explanations for this pattern in Section 5.

The Gap measure reflects the main effect of the Observable treatment—sellers after the first are strongly clustered in their sell time around the time of the first sale. 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.

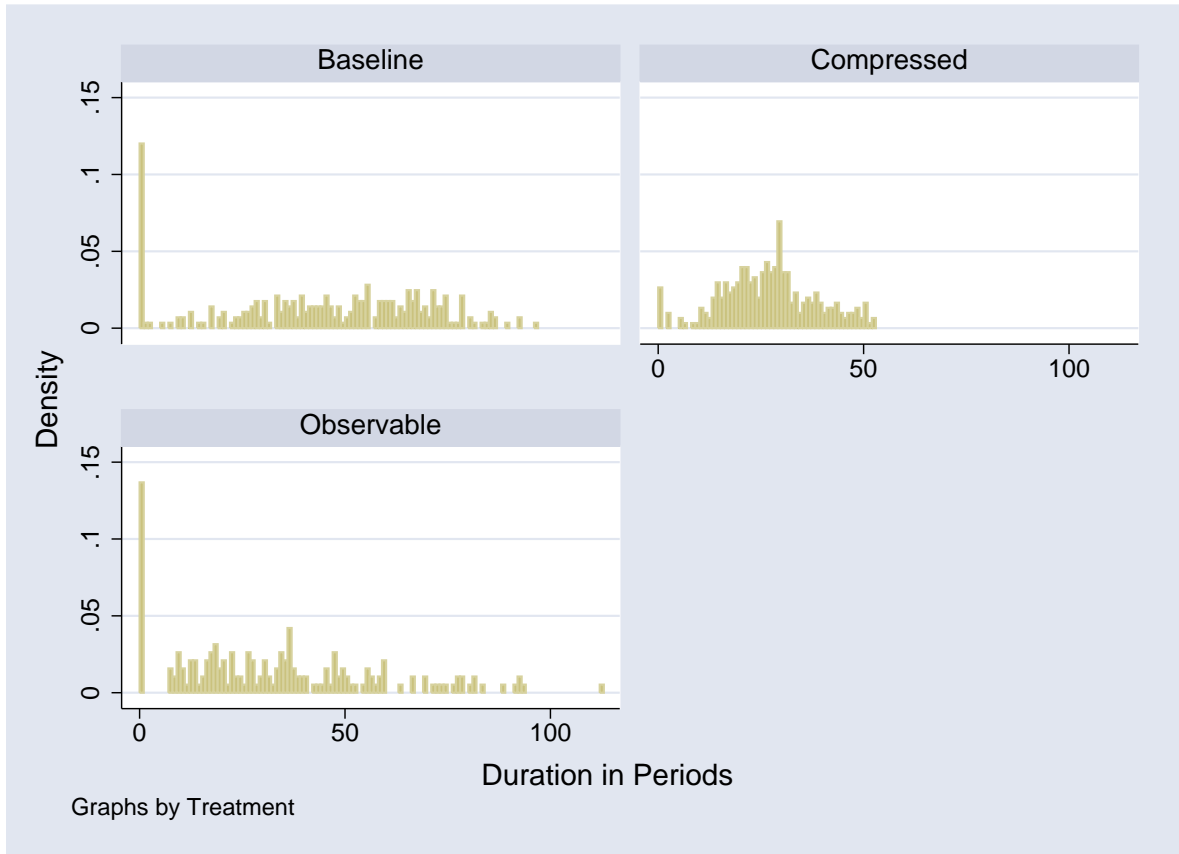


Figure 4: Duration by Treatment

To see the variation in the data within each session, we present histograms of Duration, Delay (for seller 3), and Gap (between sellers 2 and 3) in Figures 4-6. The results are qualitatively similar for other sellers. In Figure 4, 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. In Section 5, we suggest explanations for these observations based on trembles and cognitive hierarchies.

In Figure 5, histograms of Delay for the third seller are displayed for the Baseline and Compressed treatments. Since the third seller is pivotal, that subject’s Delay behavior merits particular attention. As Figure 5 shows, conditional on positive strategic delay, Delay tends to be shorter and more concentrated in Compressed sessions compared to Baseline sessions. However, third sellers are much more likely to exit before receiving a signal in Baseline compared to Compressed.

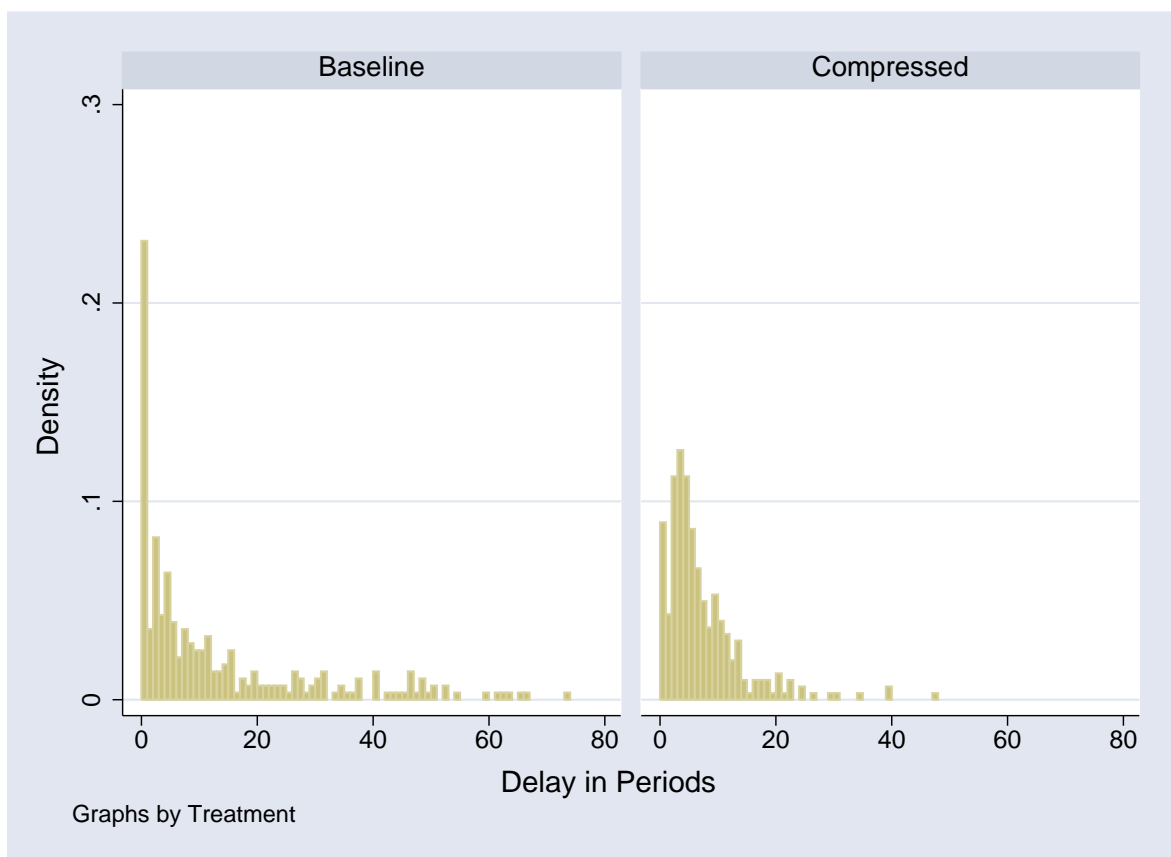


Figure 5: Third Seller Delay by Treatment

Finally, Figure 6 highlights the strong clustering effect of sales in the Observable treatment compared to the other two treatments. As the figure shows, the Gap between the second and third seller under the Observable treatment is much lower compared to the other treatments.

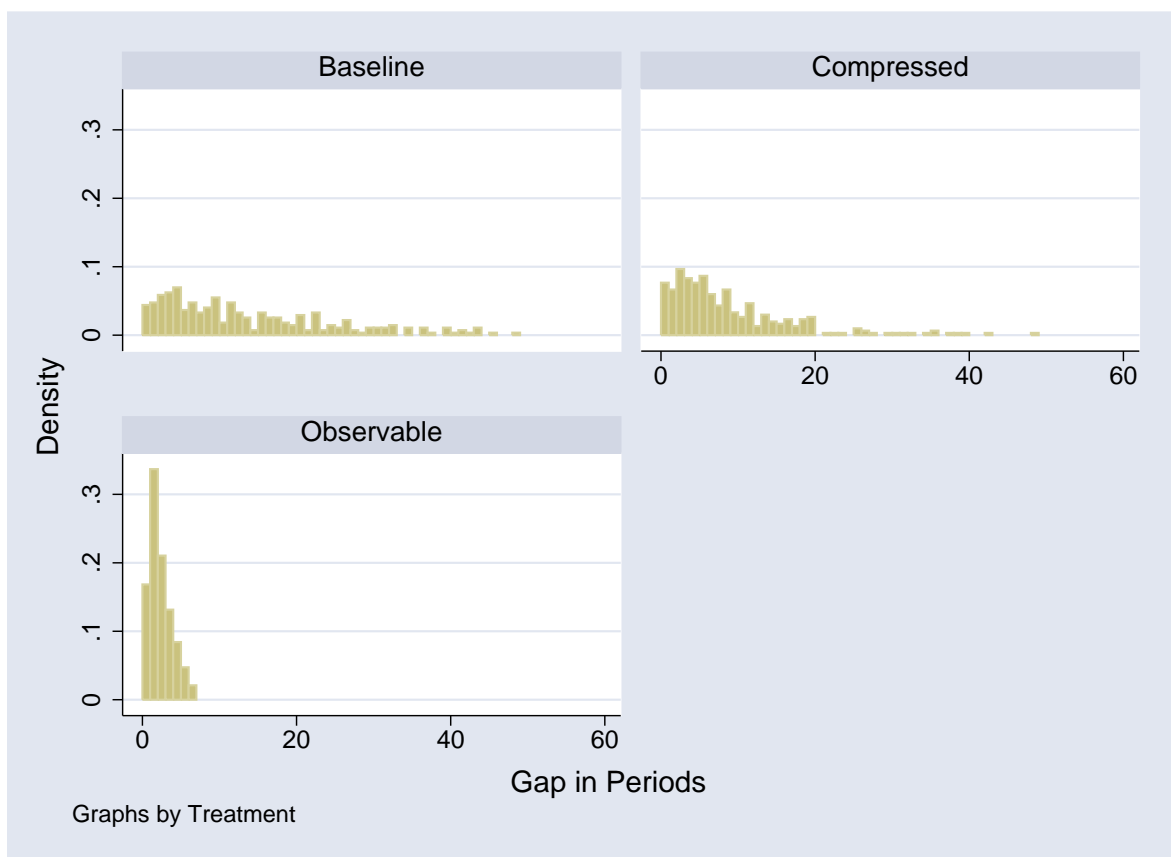


Figure 6: Gap between Seller 2 and 3 by Treatment

Taken together, there is considerable evidence for a number of treatment effects predicted by the theory; however, as the figures show, there is considerable variability in outcomes and some notable discrepancies between the theory predictions and observed behavior. Thus, our results thus far are merely suggestive of the possibility of significant treatment effects, and lots of open questions remain. In the rest of this section, we perform a variety of statistical tests to understand choice behavior and treatment effects in more detail. In Section 5 we offer some tentative explanations to try to reconcile certain discrepancies between the theory and observed behavior

4.2 Session-Level Analysis

Since all subjects interacted with one another during a session, it may not be reasonable to regard performance measures associated with an individual subject in a given round of the experiment as being statistically independent. On the other hand, since no subject participated in more than one session, it does seem reasonable to regard a session as an independent observation. Thus, in this section, we study treatment effects treating the session as the unit of observation. Obviously, this is a conservative

approach to the data—it reduces the dataset to 16 observations—nonetheless, we begin with this approach in assessing treatment effects. 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.

Relying on Propositions 2 and 4 we test the following predictions.

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 2.a *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). Note, however, that Prediction 2b is ambiguous in general. We know that the relationship described in Prediction 2b reverses when the window of awareness shrinks to fewer than 60 periods.

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 largely supported by the data—even treating the session as the unit of observation. Next, we take a less conservative approach to the data and use econometric techniques to estimate the choice strategies of individual subjects and compare these to the theory.

4.3 Individual-Level Analysis

In this section we focus on individual-level strategies. While we now treat each subject/round as a separate data point, where possible we make corrections to allow for the possibility of heteroskedasticity and autocorrelation in subjects' choices.

4.3.1 Delay

We first examine Delay. Propositions 1 and 4 offer precise predictions for individual Delay under each treatment. Specifically, the theory suggests that Delay is a constant number of periods following the signal for each seller in the case of the Baseline and Compressed treatments, and for the first seller in the case of the Observable treatment.

To examine this, we use the following regression specification:

$$DELAY_{ir} = \beta_0 + (TREATMENT_i \times t_{0,ir})\beta + \nu_r + \varepsilon_{ir}, \quad (7)$$

where i denotes the unique identifier for each subject and r denotes the round of the game. The explanatory variables are the treatments (denoted as $TREATMENT_i$, which are dummy variables for each treatment) and $t_{0,ir}$, the realization of t_0 for subject i in round r . In the tables below, we report individual and interaction effects of these variables. While the theory predicts that Delay is independent of t_0 , as we shall see, the timing of this event does play an important role in the decision to delay. To account for learning over the course of the experiment, a fixed effect ν_r for the round of the game in which the observation occurred is included. Since Delay is only meaningful for seller 1 in the Observable treatment, we run the regression specification in equation (7) separately for each seller.

By pooling the treatments and adding treatment dummies instead of regressing each treatment separately, we make better use of the data, but we implicitly assume that the error distribution is the same for all treatments. Finally, to correct for heteroskedasticity as well as correlation in the choices of a particular subject, we use robust (White-corrected) standard errors treating each subject as a “group” in constructing the variance-covariance matrix. Later, for Tobit regressions we regress each treatment separately, since the error structure has a larger impact on the estimates.

Table 3: Delay Estimates

	Robust-Cluster-OLS			Tobit	
	Seller 1	Seller 2	Seller 3	Baseline	Compressed
Constant	12.841 (10.78)**	17.931 (10.30)**	22.803 (10.28)**	17.277 (5.31)**	11.011 (7.46)**
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)**		
$t_0 \times \text{Compressed}$	0.045 (5.03)**	0.064 (4.59)**	0.086 (4.82)**		
$t_0 \times \text{Observable}$	0.012 (0.93)				
Round Fixed Effects	Yes	Yes	Yes	Yes	Yes
Observations	738	584	583	1681	1788
<i>R</i> -squared	0.23	0.26	0.28		

OLS: Robust t -statistics in parentheses. Tobit: Standard t -statistics in parentheses
 * significant at 5%; ** significant at 1%

OLS Regressions. The first column of Table 3 shows the results of this analysis restricting attention to subjects who execute the first sale in each round of the experiment.¹⁶ The regression coefficient estimates imply that the Compressed treatment reduces Delay by 3.48 periods, which is statistically different from zero at the 1% significance level (F -statistic 19.25), but far from the 18 period reduction predicted by the theory.¹⁷ The regression coefficient estimates imply that the Observable treatment reduces Delay by 1.27 periods, although this is statistically indistinguishable from zero at conventional significance levels (F -statistic 1.14). These results are in the direction predicted by the theory, but clearly inconsistent with the level predictions.

¹⁶In the event that multiple sellers sold in the same period and no sales were executed prior to this period, we randomly assign one of these sellers the identity of “first” seller.

¹⁷Recall that the estimated mean marginal effect of the Compressed treatment is equal to the Compressed coefficient plus the coefficient of the interaction term times the sample mean of t_0 , which was approximately equal to 75 in the dataset.

Columns two and three of Table 3 report analogous results for the second and third subject to sell in a given round of the experiment. Since under the Observable treatment Delay is not a meaningful measure for sellers beyond the first, we omit observations under this treatment. As Table 3 shows, the results for the second and third seller are qualitatively similar with what was observed looking only at the first seller. One interesting difference is that second and third sellers tend to delay longer than the first seller under both the Baseline and Compressed treatments. The point estimates for the Delay reduction associated with the Compressed treatment are 5.39 periods (for the second seller) and 6.45 periods (for the third seller). Both point estimates are statistically different from zero, but also significantly different from the theory prediction as well.

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. 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. It is, however interesting to notice that the t_0 effect is systematically less pronounced in the Compressed treatment compared to Baseline or Observable. We return to this in more detail in the next section.

Tobit Regressions. Note, however, that the OLS regressions restrict attention only to observations where a sale takes place. Thus, our data is censored. Only for subjects who received the signal and successfully exited the game, do we directly observe Delay. Thus, we are potentially omitting a considerable amount of information on individual choices. For example, subjects who were planning to exit (sell) in period T_i^* have been excluded from the analysis up to now if $T_i^* \geq T^{\text{end-of-game}}$. We know only that their exit point $T_i^* \geq T^{\text{end-of-game}}$; thus the Delay for these subjects is right censored. Using the Tobit estimation procedure, in principle, allows us to address the right censoring problem and use more of the data.

Specifically, in the Baseline and Compressed treatments, player i 's optimal strategy is to exit at $T_i^* = t_i + \tau$. Since players err, we assume

$$T_i^* = t_i + \tau + \varepsilon_i,$$

where $\varepsilon_i \sim \mathcal{N}(0, \sigma_{\text{treatment}}^2)$. 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 we run separate estimations for each treatment.¹⁸ The theory model then implies that we can make the following transformation

$$DELAY_i^* = T_i^* - t_i = \tau + \varepsilon_{ir},$$

¹⁸Of 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 $DELAY_i^*|t_i \sim \mathcal{N}(\tau, \sigma_{\text{treatment}}^2)$. Using this transformation, we are now in a position to use the Tobit procedure to obtain estimates of Delay under the Baseline and Compressed treatments.¹⁹ These estimates are reported in columns 4 and 5 of Table 3. Notice that the estimated Delay for the Baseline treatment (17.277) is considerably higher than the Delays for sellers 1 through 3 shown in Table 2.

Turning to the Compressed treatment, notice that the estimated Delay is considerably lower (11.011) than under the Baseline treatment as predicted by the theory. Formally, one can easily reject the null hypothesis that the parameter estimate for Delay length are equal at the 1% significance level. In contrast to Baseline, the estimated Delay under the Compressed treatment is *larger* than the theory prediction (Delay = 5), and one can reject the null hypothesis implied by the theory against the two-sided alternative at the 1% significance level.

To summarize, it is reassuring that the two main predictions of the theory—that there is significant strategic delay and that delay is increasing in the length of the window of awareness are borne out in both the OLS and Tobit estimates. That being said, a key difficulty with the Tobit estimation procedure is that it (of necessity) excludes factors such as the time that the true value of the asset stopped growing (t_0) as well as round fixed effects that we know, from the OLS regressions, do affect subject choices. Thus, our view is the main value of this analysis is in capturing qualitative treatment effects rather than in generating exact point estimates. Indeed, taken as a whole, Table 3 shows that point estimates of Delay can vary a good deal depending on the sample and estimation procedure employed.

4.3.2 Herding Behavior

Next, we turn to estimates of herding behavior by subjects. In the beginning of this section we introduced the measure GAP^{kl} , which measures the number of periods between the sale times of the l th seller and the k th seller. Recall that the theory predicts that Gap is shorter for the Compressed than for the Baseline treatment. For the Observable treatment, the theory predicts zero Gaps. Note however that, owing to the discretization of time in the experimental implementation of the theory, information about time of the first sale is delayed by about a half second (one period); hence, the theory effectively implies that $GAP^{12} = 1$ and $GAP^{23} = 0$ in the Observable treatment. To examine this, we run the following regression:

$$GAP_{ir}^{kl} = \beta (TREATMENT_i \times t_{0,ir}) + \nu_r + \varepsilon_{ir}, \quad (8)$$

where the variables are defined as in equation (7).

¹⁹As for the case with the second and third sellers in the OLS regressions, the analysis is not meaningful in the case of the Observable treatment.

Table 4: Gap Estimates

	Gap 1 st & 2 nd	Gap 2 nd & 3 rd
Constant	15.631 (8.66)**	9.942 (7.21)**
Compressed	-9.061 (3.55)**	-5.67 (3.45)**
Observable	-10.057 (4.73)**	-8.149 (5.51)**
t₀	0.095 (3.11)**	0.066 (3.55)**
t₀ × Compressed	0.069 (1.61)	-0.005 (0.22)
t₀ × Observable	-0.114 (3.42)**	-0.066 (3.27)**
Round Fixed Effects	Yes	Yes
Observations	774	775
<i>R</i> -squared	0.20	0.29

Robust *t*-statistics in parentheses

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

The first column of Table 4 shows the results of this analysis comparing the 1st and 2nd sellers. The main interest of this analysis is the effect of the Observable treatment on the Gap measure. The theory predicts that when selling is observable, there will be substantial herding following the first sale. This effect is borne out in the data. Compared to Baseline, the coefficient estimates in the first column of Table 4 indicate that the Observable treatment reduces the Gap by 18.61 periods. This is significantly different from zero at the 1% significance level (*F*-statistic 92.38). The theory predicts a Gap reduction of 12 periods in going from Baseline to Observable. However, as Table 2 shows, for the Baseline treatment, the Gap is 23 periods. Thus, the “herding” hypothesis would imply a reduction of 22 periods in going from Baseline to Observable. Indeed, we cannot reject the null hypothesis of a 22 period Gap reduction (*F*-statistic 2.98). The second column of Table 4 examines the Gap between the 2nd and 3rd sellers. Here, the regression estimates imply a Gap reduction of 13.10 periods in going from Baseline to Observable. This is extremely close to the theoretical prediction of herding. Indeed, we cannot reject the herding hypothesis at conventional significance levels. Taken together, we find strong evidence in support of the herding prediction.

Next, turning to the Compressed treatment, notice that the regression estimates in column 1 of Table 4 imply a 3.88 period reduction in the Gap between the 1st and 2nd sellers in going from Baseline to Compressed. While this is in the direction predicted by the theory, we cannot reject the null hypothesis of no treatment effect at conventional significance levels (F -statistic 2.40). The second column of Table 4 shows that the Gap between the 2nd and 3rd sellers is reduced by 6.05 periods in going from Baseline to Compressed. Here we can reject the null hypothesis of no treatment effect in favor of the one-sided alternative predicted by the theory at the 1% significance level (F -statistic 39.01). The exact theory prediction is a Gap reduction of 6 periods.²⁰ Indeed, we cannot reject the theory prediction at any level (F -statistic 0.00). Thus, there is support for the theory prediction of shorter gaps with shorter windows of awareness.

Finally, notice that there is a significant effect on Gaps associated with t_0 —the later is t_0 , the longer the gap between sales. In contrast to Table 3, the interaction effect of t_0 with the Compressed treatment is no longer significant while the interaction of t_0 with Observable becomes highly significant. The theory, of course, predicts that all of these coefficients should be zero.

5 Discussion

As the above analysis shows, the theory model does well at predicting several aspects of the data. However, there are a number of puzzling discrepancies between the theory and actual behavior. In this section, we highlight a number of these discrepancies and offer some post hoc rationalization for what might be going on.

Perhaps the central prediction of the theory model is that players should delay exiting until after receiving their signal or, in the case of the Observable treatment, after observing the time of the first exit. Yet, as Figure 5 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 left-hand side variable, *EARLY-EXIT* equals one if a subject sold (weakly) prior to receiving the signal and zero otherwise. Our regressors are t_0 , the treatments, and the interaction between these variables, and are reported in Table 5 below. Column 1 of Table 5 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 examines the results, restricting the sample to the first seller only for all treatments.

²⁰To obtain this number we take difference between expected time that the third trader receives the signal for the two different values of η . That is, $\frac{1}{7}90 - \frac{1}{7}50$.

Table 5: Probit Model of Probability of *EARLY-EXIT*

	Baseline and Compressed Only	Seller 1 Only
Constant	−2.865 (12.44)**	−1.892 (4.24)**
Compressed	0.18 (0.75)	0.469 (1.41)
Observable		0.697 (2.18)*
t_0	0.02 (14.77)**	0.022 (8.64)**
$t_0 \times$ Compressed	−0.006 (3.54)**	−0.008 (2.52)*
$t_0 \times$ Observable		−0.009 (2.69)**
Round Fixed Effects	Yes	Yes
Observations	2259	738

Robust z -statistics in parentheses

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

Turning to Column 1, the Probit analysis estimates a probability of early exit (i.e. selling before receiving a signal) in the Baseline treatment of 9.04%. The marginal effect of the Compressed treatment is to *reduce* the probability of early exit. The marginal effect of an increase in t_0 is to increase the probability of early exit by 0.3% per period. Restricting attention only to the first seller (Column 2, Table 5), the probability of early exit in the Baseline treatment increases to 22.22%. The marginal effect of the Compressed and Observable treatments is to modestly reduce the probability of early exit compared to Baseline. Finally, increases in t_0 lead to increases in early exit; in this case, the marginal effect is 0.66% per period.

Clearly, the significant amount of early exit behavior and its dependence on both the treatment and the timing of the t_0 event is at odds with the theory. To see the impact of this behavior on outcomes, it is useful to consider the probability of a zero Duration event. In a zero Duration event, three or more subjects exit prior to the t_0 event. The theory model predicts that there should be no zero duration events; however, in our sessions, zero Duration events occurred 15% of the time in the Baseline and Observable treatments but less than 5% of the time in Compressed.

What accounts for the differences in early exit behavior in the Compressed and Observable treatments compared Baseline? Why does the probability of early exit depend on the time period of the t_0 event? Why are zero duration events more common in Baseline and Observable compared to the Compressed treatment?

To rationalize the finding that the probability of early exit depends on the timing of the t_0 event, suppose that subjects use the following heuristic strategy: A subject tries to anticipate how far in the past (if at all) the t_0 event has occurred. Once the expected time since the t_0 event has occurred grows sufficiently large, a subject will exit. Suppose a subject receives a signal at time t . In that case, the expected time since the t_0 event has occurred is the same regardless of t owing to the exponential process generating t_0 . Thus, a subject following the heuristic strategy would simply wait a fixed number of periods after having received the signal before exiting. Notice that the equilibrium strategy derived earlier is a special case of this heuristic. This perhaps explains the performance of the equilibrium model in predicting the main treatment effects we observe despite the apparent complexity of the experimental setting.

Next, suppose that subjects following this heuristic strategy suffer from the following cognitive bias: Subjects perceive an increasing arrival rate of the t_0 event over time instead of the actual constant arrival rate. Note that the vast majority of experiences individuals have with random arrival processes have increasing, rather than constant, arrival rates. If subjects perceive the arrival rate as increasing, then, as the time period of the game increases, subjects not receiving a signal think it increasingly likely that the t_0 event has occurred but they are uninformed about it. Once this becomes sufficiently likely, a subject following the heuristic strategy described above will (rationally) choose to exit rather than to stay in the game. Therefore, the incidence of early exits should be increasing in t_0 —across all treatments. This is consistent with the results in Table 5 as well as on the coefficient associated with t_0 in explaining Delay in Table 3.

Bias in the perception of arrival rates can also explain differences across treatments. Notice that, in the Compressed treatment, subjects realize that the “window of awareness” is shorter and, therefore, correctly perceive that, even if they have not received the signal, the t_0 event could not have occurred too far in the past. Hence, the incentives to exit early are reduced in this treatment relative to Baseline. In the Observable treatment, the expectation about the timing of the t_0 event is identical to Baseline; however the potential downside from remaining in the game is lower in Observable since a subject who does not exit early still has a 2 in 5 chance of exiting successfully after observing the first exit.

How does this rationale compare with alternative explanations? A simple alternative is that players simply make random errors in their exit times. As we will show below, the simplest version of this explanation is unsatisfactory in explaining zero duration events in the data. To see this, first consider the Observable treatment. In this treatment, such errors can give rise to herding behavior and hence to zero duration events. Formally, suppose that there is a chance ε that each seller sells before t_0 and

that all others herd immediately after the first sale. Thus, to obtain zero duration events the observed 13.68% of the time in this model requires that ε solve

$$1 - (1 - \varepsilon)^6 = .1368$$

which yields an error rate of 2.4%. This does not seem unreasonable given the complexity of the experiment. However, applying this same error rate to the Baseline and Compressed treatments does not yield sensible results. Specifically, this same 2.4% error rate leads to the prediction that, in both the Baseline or Compressed treatments, zero duration events occur with probability

$$\sum_{j=3}^6 \binom{6}{j} \varepsilon^j (1 - \varepsilon)^{6-j} |_{\varepsilon=.024} = .0003$$

This is a poor prediction in terms of the levels of zero duration events as well as for differences across treatments. (Recall that zero duration events occur approximately three times as often in Baseline as in Compressed, yet the error explanation predicts no difference.)

The error explanation is, however, quite useful at explaining the relationship between the timing of the t_0 event and the Gap between the first and second sales. Recall from Table 4 that in the Baseline and Compressed treatment, the Gap between the first and second sales is increasing in t_0 ; whereas in the Observable treatment, it is decreasing in t_0 .

Notice however, that the error rationale described above mostly ignores the strategic effects of the possibility of errors on equilibrium strategies. If others are making mistakes, this changes the cost/benefit trade-off for a subject and hence has strategic effects and it is possible that these strategic effects might help to rationalize the discrepancies between the data and the theory model. One avenue to formalize these ideas is to study clock games using the quantal response equilibrium solution concept (see, for example, McKelvey and Palfrey (1995) and Capra, Goeree, Gomez, and Holt (1999)). While this seems promising, the complexity of equilibrium calculation under the assumption of full rationality led us to leave this analysis for future research.

6 Extending Clock Games: Effects of Information Clustering

There are two key differences between the “standard” model of clock games we just presented and its nearest antecedent, the AB model. First, in the standard model there are a discrete number of players whereas in AB there is a continuum. Second, in the standard model information is clustered, in the sense that each player possesses a positive fraction of the aggregate information available about t_0 . In AB, however,

information is (uniformly) diffused and each player possesses a zero measure of the aggregate information. In many situations, it seems more realistic to that there are only a few key players and/or their information is clustered. This is clearly the case in the situation of new product introductions. Similarly, in the situation of a currency attack, there will only be a few key institutions with the information and the financial resources to move effectively against the central bank. Next, consider the case of the stock market. Here there are many individuals but information is arguably clustered. For instance, suppose that a “cohort” of investors has as their primary news source *The Wall Street Journal* while others rely mainly on *Nightly Business Review*. The difference in news coverage between these two sources naturally leads to information clustering.

To isolate the pure effect of information clustering, we extend the basic clock games model to study a hybrid model consisting of a continuum of individuals but where information clustering is present. We refer to it as the CC (continuum with clustering) model. In this model, players occur on a continuum but are divided into I equal sized “cohorts.” All players in a cohort receive the signal at an identical point in time—so information is clustered. The game ends when a mass of players equal to K cohorts decides to exit. We then study the relation between equilibrium delay in the extended model with the two polar cases—the AB model and the standard Clock Games model—to illustrate the economic effect of information clustering.

Unobservable Actions. We begin by supposing that exit decisions are unobservable and characterizing the set of symmetric equilibria in the CC model. It is useful to define:

$$\tau(K, I) = \frac{1}{g} \ln \left(\frac{\lambda F(K, I, \eta(\lambda - g))}{g - (g - \lambda) F(K, I, \eta\lambda)} \right).$$

This is identical to the expression for τ given in Proposition 1 but with functional arguments for K and I added.

Proposition 5 *In the CC model, any delay $\tau_{CC} \in [\tau(K - 1, I), \tau(K, I)]$ is a symmetric equilibrium.*

Notice that the equilibrium in Proposition 5 exhibiting the longest equilibrium delay corresponds exactly to the unique symmetric equilibrium in the standard model. Moreover, this equilibrium is the Pareto-best of all the equilibria arising in the CC model.

What accounts for the equilibrium multiplicity when there is a continuum of players compared to the standard case? The key is that the presence of information clustering with a continuum of players introduces a discontinuity in the expected payoffs of player i in a given cohort. To see this, notice that if player i exits just a bit before his cohort, he faces a trade-off between the gains from waiting versus the chance that the game

will end—this is just the chance that the K th of $I - 1$ other cohorts will exit. On the other hand, by waiting just a bit after his own cohort exits, player i again faces a trade-off between the gain from waiting and the loss due to the chance that the game will end—however this latter event now occurs when the $K - 1$ th of the $I - 1$ other cohorts exits. Thus, the risk of the game ending is discretely higher for player i just after his own cohort exits compared to just before. One way to see this is to examine the marginal cost curve (expressed in growth rates) faced by player i if his cohort is set to exit at time $t_i + \tau$. This is illustrated in Figure 3 below for the parameter values used in the experiment.




Figure 3: Marginal benefit, g , versus marginal cost, MC , for different τ

The horizontal line at $g = 0.02$ reflects the marginal benefit from waiting. The lowermost marginal cost curve reflects the chance that K out of $I - 1$ remaining cohorts will exit. Where this line crosses the marginal benefit line corresponds to the longest equilibrium delay, $\tau(K, I)$. The uppermost marginal cost curve reflects the chance that $K - 1$ out of $I - 1$ remaining cohorts will exit at a given time. Where this line crosses the marginal benefit line corresponds to the shortest equilibrium delay, $\tau(K - 1, I)$. Because information is clustered, player i knows that his cohort exits exactly at time $t_i + \tau_{CC}$. Thus, if player i waits until after this time, the relevant marginal cost he faces is the uppermost marginal cost curve. Similarly, if player i considers exiting just before time $t_i + \tau_{CC}$, the relevant marginal cost curve he faces corresponds to the lowermost marginal cost curve. Thus, player i faces a discontinuity in his marginal cost

at time $t_i + \tau$, and this discontinuity in turn leads to equilibrium multiplicity. By way of comparison, no such discontinuity arises in the standard model in which information clustering is also present. Interestingly, since $\tau_{CC} \leq \tau(K, I)$, player i never has an incentive to preempt his own cohort.

How does equilibrium delay in the CC model compare to that in the AB model? In the next proposition, we show that information clustering can lead to strategic incentives for increased equilibrium delay. For future reference, it is useful to denote the equilibrium delay arising in the AB model where $\kappa = \frac{K}{I}$ as:

$$\tau_{AB} = \frac{1}{g} \ln \left(\frac{\lambda e^{\kappa\eta(\lambda-g)}}{g - (g - \lambda) e^{\kappa\eta\lambda}} \right).$$

Proposition 6 *The longest equilibrium delay in the CC model exceeds that without information clustering. Formally, $\tau(K, I) > \tau_{AB}$.*

An immediate corollary of this Proposition is:

Corollary 1 *Equilibrium delay in the standard model with unobservable actions always exceeds equilibrium delay in the AB model.*

The relationship between equilibrium delay in the standard model and that in AB is again readily illustrated in Figure 3. First, notice that the marginal benefit of delay is the same in the two models. The marginal cost in the AB model is illustrated by the middle marginal cost line in the figure—this always lies above the marginal cost curve associated with the chance that the K th of $I - 1$ cohorts will exit—that is, the marginal cost curve in the standard model. The figure thus illustrates how information clustering leads to incentives for longer equilibrium delay.

A different way to see this same effect is to consider the case where the number of cohorts grows arbitrarily large. Clearly, in this case, information clustering becomes vanishingly small. We show in the next proposition that this in turn leads to convergence in the set of equilibria in the CC model—as the number of cohorts increases arbitrarily, all equilibria converge to that of the AB model.

Proposition 7 *As the number of cohorts grows arbitrarily large, the set of symmetric equilibria in the CC model converges to the equilibrium delay in the model without information clustering.*

Formally, $\lim_{I \rightarrow \infty} \tau(\kappa I - 1, I) = \lim_{I \rightarrow \infty} \tau(\kappa I, I) = \tau_{AB}$.

Observable Actions. Next, we turn to the CC model when actions are observable. Analogous to Proposition 5, we have

Proposition 8 *In the CC model with observable actions, the following comprises a symmetric equilibrium: After any player has exited all other players exit immediately. Prior to any player exiting all players delay for any $\tau_{1,CC} \in [0, \tau_1]$.*

As in the previous case, the longest equilibrium delay in the CC model with observable actions is identical to the unique symmetric equilibrium in the standard model identified in Proposition 4. Unlike the standard model, any shorter delay is also consistent with equilibrium. Why is this? The key is that, with a continuum of players, having received his signal it is not possible for player i to obtain any benefit from waiting after his cohort has exited since the game ends immediately. Since player i 's marginal benefit always exceeds his marginal cost for delay times shorter than $\tau_{1,CC}$, there is again no benefit to i from preempting his own cohort. Thus, player i again faces a discontinuity in his marginal cost of waiting at the time that his own cohort exits, which is absent in the standard model.

How does this compare with the AB model when actions are observable? While this case is not analyzed in AB (2002, 2003), we derive here the following result.

Proposition 9 *In the AB model with observable actions, the unique symmetric equilibrium is for the first player to exit immediately upon receiving the signal, i.e. $\tau_{1,AB} = 0$, and for all remaining players to exit immediately thereafter.*

Recall that equilibrium delay in the standard model (Proposition 4) occurs since player i is uncertain when the next player will exit. In a model without information clustering, the exact degree of synchronization is common knowledge. Hence, each player knows that an instant after t_i the next player will receive a signal and exit. In short, the exact knowledge of the degree of synchronization destroys equilibrium delay. A different way to see this is to consider the case in the CC model where the number of cohorts grows arbitrarily large. In the limit, the exact degree of synchronization again becomes known to all players and, as we show below, there is no equilibrium delay in the limit.

Proposition 10 *In the CC model with observable actions, there is no equilibrium delay as the number of cohorts grows arbitrarily large. Formally, $\lim_{I \rightarrow \infty} \tau_1 = \tau_{1,AB} = 0$.*

7 Conclusions

We introduced and analyzed a class of games we refer to as clock games. These games show how equilibrium delay caused by the economic trade-off between the waiting and preemption motives first introduced in the Abreu and Brunnermeier (2003) (AB) model, can be extended to a broad class of economic situations. Importantly, unlike the AB model, our more general clock games framework is amenable to empirical testing. In addition, we showed that the Clock Games model differ from the AB model in its predictions about the effect of transparency (observability of other players' moves) on equilibrium delay—the Clock Games framework predicts delay followed by herding while the AB model predicts no delay.

We then tested the model using controlled laboratory experiments. Our key treatments were to vary the observability of moves and to vary 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 we found support for many of the qualitative predictions of the clock games model, 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.
- 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.
- CAPRA, C. M., J. K. GOEREE, R. GOMEZ, AND C. A. HOLT (1999): “Anomalous Behavior in a Traveler’s Dilemma,” *American Economic Review*, 89(3), 678–690.
- 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.
- 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.

- McKELVEY, R. D., AND T. R. PALFREY (1992): "An Experimental Study of the Centipede Game," *Econometrica*, 60(4), 803–836.
- (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10(1), 6–38.
- 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

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

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 F 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

Suppose that all other players follow τ strategies. We establish that player i cannot profitably deviate by exiting after $\tau + \Delta$ periods for any $\Delta \neq 0$. For convenience, we shall refer to such a strategy as a “ Δ strategy” while a strategy where $\Delta = 0$ we shall call a “ τ strategy.” Let $z := t_i - t_0$. We first derive the hazard rate and expected drop in payoff when i follows a Δ strategy given that all other players are using τ strategies.

A.2.1 Deriving the Hazard Rate

Consider the situation of player i . In an equilibrium in τ strategies, the probability that the game ends at $t_i + \Delta + \tau$ is equal to the probability that the K^{th} of the other $(I - 1)$ players received a signal at $t_i + \Delta$. That is, the probability that (i) $(K - 1)$ of the other players receive a signal prior to $t_i + \Delta$ (ii) one player receives a signal at $t_i + \Delta$ and (iii) the remaining $(I - 1 - K)$ players receive signals after $t_i + \Delta$. Formally, let $\Pi(\Delta|t_0)$ denote the probability that player j receives a signal prior to $t_i + \Delta$ for a given t_0 . That is, $\Pi(\Delta|t_0) \equiv \Pr(x \leq t_i + \Delta|t_0) = \frac{t_i + \Delta - t_0}{\eta}$, and let $\pi(\Delta|t_0)$ denote the corresponding density. Let D_Δ denote the event that the game ends at $t_i + \Delta + \tau$ and let $f(\cdot|t_0)$ denote the density of this event conditional on t_0 . For $\Delta \geq 0$ and $t_0 \in [t_i + \Delta - \eta, t_i]$

$$\begin{aligned} f(D_\Delta|t_0) &= \binom{I-1}{1} \pi(\Delta|t_0) \binom{I-2}{K-1} [\Pi(\Delta|t_0)]^{K-1} [1 - \Pi(\Delta|t_0)]^{I-1-K} \\ &= \frac{(I-1)!}{(K-1)!(I-1-K)!} \left(\frac{1}{\eta}\right)^{I-1} (t_i + \Delta - t_0)^{K-1} (\eta - (t_i + \Delta - t_0))^{I-1-K}. \end{aligned}$$

More generally, using $z := t_i - t_0$, we have

$$f(D_\Delta|t_0) = \begin{cases} 0 & \text{for } t_i + \Delta < t_0 \\ \frac{(I-1)!}{(K-1)!(I-1-K)!} \left(\frac{1}{\eta}\right)^{I-1} (\Delta + z)^{K-1} (\eta - (\Delta + z))^{I-1-K} & \text{for } t_i + \Delta \in [t_0, t_0 + \eta] \\ 0 & \text{for } t_i + \Delta > t_0 + \eta. \end{cases}$$

Suppose that at time $t_i + \Delta + \tau$, i knows that the game is still ongoing. This means that, at most, $(K-1)$ of the other players received a signal prior to $t_i + \Delta$. Let us denote the event that the game is still “alive” at $t_i + \Delta + \tau$ by A_Δ^- .

For $\Delta \geq 0$ and $t_i + \Delta \in [t_0, t_0 + \eta]$:

$$\begin{aligned} \Pr(A_\Delta^-|t_0) &= \sum_{n=0}^{K-1} \binom{I-1}{n} \Pi(\Delta|t_0)^n [1 - \Pi(\Delta|t_0)]^{I-1-n} \\ &= \sum_{n=0}^{K-1} \frac{(I-1)!}{n!(I-1-n)!} \left(\frac{1}{\eta}\right)^{I-1} (t_i + \Delta - t_0)^n (\eta - (t_i + \Delta - t_0))^{I-1-n}. \end{aligned}$$

More generally, using $z := t_i - t_0$, we have

$$\Pr(A_\Delta^-|t_0) = \begin{cases} 1 & \text{for } t_i + \Delta < t_0 \\ \sum_{n=0}^{K-1} \frac{(I-1)!}{n!(I-1-n)!} \left(\frac{1}{\eta}\right)^{I-1} (z + \Delta)^n (\eta - (z + \Delta))^{I-1-n} & \text{for } t_i + \Delta \in [t_0, t_0 + \eta] \\ 0 & \text{for } t_i + \Delta > t_0 + \eta \end{cases}$$

However, player i does not know t_0 . Hence, we have to take expectations over all possible t_0 . Recall that t_0 is exponentially distributed, with p.d.f.

$$\phi(t_0) = \lambda e^{-\lambda t_0}.$$

Further, i 's signal, t_i , implies that $t_i - \eta \leq t_0 \leq t_i$, hence

$$\begin{aligned} \phi(t_0|t_i) &= \frac{\lambda e^{-\lambda t_0}}{1 - e^{-\lambda t_i} - (1 - e^{-\lambda(t_i - \eta)})} \\ &= \frac{\lambda e^{\lambda z}}{e^{\lambda \eta} - 1}. \end{aligned}$$

Thus, player i 's hazard rate of the game ending at time $t_i + \Delta + \tau$ is

$$\begin{aligned} h(\Delta, t_i) &= \frac{f(D_\Delta|t_i)}{\Pr(A_\Delta^-|t_i)} \\ &= \frac{\int_{t_i - \eta}^{t_i} \phi(t_0|t_i) f(D_\Delta|t_0) dt_0}{\int_{t_i - \eta}^{t_i} \phi(t_0|t_i) \Pr(A_\Delta^-|t_0) dt_0}. \end{aligned}$$

When $\Delta \geq 0$, this expression reduces to

$$\begin{aligned}
h(\Delta, t_i) &= \frac{\int_{t_i+\Delta-\eta}^{t_i} \frac{\lambda e^{\lambda(t_i-t_0)}}{e^{\lambda\eta-1}} \frac{(I-1)!}{(K-1)!(I-1-K)!} \left(\frac{1}{\eta}\right)^{I-1} (t_i + \Delta - t_0)^{K-1} (\eta - (t_i + \Delta - t_0))^{I-1-K} dt_0}{\int_{t_i+\Delta-\eta}^{t_i} \frac{\lambda e^{\lambda(t_i-t_0)}}{e^{\lambda\eta-1}} \sum_{n=0}^{K-1} \frac{(I-1)!}{n!(I-1-n)!} \left(\frac{1}{\eta}\right)^{I-1} (t_i + \Delta - t_0)^n (\eta - (t_i + \Delta - t_0))^{I-1-n} dt_0} \\
&= \frac{\int_0^{\eta-\Delta} e^{\lambda z} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z + \Delta)^{K-1} (\eta - (z + \Delta))^{I-1-K} dz}{\int_0^{\eta-\Delta} e^{\lambda z} \sum_{n=0}^{K-1} \frac{(I-1)!}{n!(I-1-n)!} (z + \Delta)^n (\eta - (z + \Delta))^{I-1-n} dz}.
\end{aligned}$$

When $\Delta < 0$, this expression reduces to

$$\begin{aligned}
h(\Delta, t_i) &= \frac{\int_{t_i-\eta}^{t_i+\Delta} \frac{\lambda e^{\lambda(t_i-t_0)}}{e^{\lambda\eta-1}} \frac{(I-1)!}{(K-1)!(I-1-K)!} \left(\frac{1}{\eta}\right)^{I-1} (t_i+\Delta-t_0)^{K-1} (\eta-(t_i+\Delta-t_0))^{I-1-K} dt_0}{\int_{t_i-\eta}^{t_i+\Delta} \frac{\lambda e^{\lambda(t_i-t_0)}}{e^{\lambda\eta-1}} \sum_{n=0}^{K-1} \frac{(I-1)!}{n!(I-1-n)!} \left(\frac{1}{\eta}\right)^{I-1} (t_i+\Delta-t_0)^n (\eta-(t_i+\Delta-t_0))^{I-1-n} dt_0 + \int_{t_i+\Delta}^{t_i} \frac{\lambda e^{\lambda(t_i-t_0)}}{e^{\lambda\eta-1}} dt_0} \\
&= \frac{\int_{-\Delta}^{\eta} e^{\lambda z} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z + \Delta)^{K-1} (\eta - (z + \Delta))^{I-1-K} dz}{\int_{-\Delta}^{\eta} e^{\lambda z} \sum_{n=0}^{K-1} \frac{(I-1)!}{n!(I-1-n)!} (z + \Delta)^n (\eta - (z + \Delta))^{I-1-n} dz + \eta^{I-1} \int_0^{-\Delta} e^{\lambda z} dz} \\
&= \frac{\int_{-\Delta}^{\eta} e^{\lambda z} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z + \Delta)^{K-1} (\eta - (z + \Delta))^{I-1-K} dz}{\int_{-\Delta}^{\eta} e^{\lambda z} \sum_{n=0}^{K-1} \frac{(I-1)!}{n!(I-1-n)!} (z + \Delta)^n (\eta - (z + \Delta))^{I-1-n} dz + \eta^{I-1} \frac{e^{-\lambda\Delta}-1}{\lambda}}.
\end{aligned}$$

A.2.2 Deriving $E[e^{-g(t_i+\Delta-t_0)} | D_\Delta, t_i]$

Recall that

$$\begin{aligned}
\phi(t_0 | D_\Delta, t_i) &= \frac{\phi(t_0 \cap D_\Delta | t_i)}{f(D_\Delta | t_i)} \\
&= \frac{f(D_\Delta | t_0) \phi(t_0 | t_i)}{\int_{t_0=t_i-\eta}^{t_i} f(D_\Delta | t_0) \phi(t_0 | t_i) dt_0}.
\end{aligned}$$

When $\Delta \geq 0$, this expression reduces to

$$\phi(t_0 | D_\Delta, t_i) = \begin{cases} \frac{e^{\lambda z} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K}}{\int_0^{\eta-\Delta} e^{\lambda z} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K} dz} & \text{for } t_0 \in [t_i + \Delta - \eta, t_i] \\ 0 & \text{otherwise} \end{cases}$$

When $\Delta < 0$, this expression reduces to

$$\phi(t_0|D_\Delta, t_i) = \begin{cases} \frac{e^{\lambda z} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K}}{\int_{-\Delta}^{\eta} e^{\lambda z} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K} dz} & \text{for } t_0 \in [t_i - \eta, t_i + \Delta] \\ 0 & \text{otherwise} \end{cases}$$

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

When $\Delta \geq 0$, this expression reduces to

$$\begin{aligned} E[e^{-g(\Delta+z)}|D_\Delta, t_i] &= \frac{\int_0^{\eta-\Delta} e^{-g(z+\Delta)} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K} e^{\lambda z} dz}{\int_0^{\eta-\Delta} e^{\lambda z} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K} dz} \\ &= \frac{\int_0^{\eta-\Delta} e^{-g(z+\Delta)} e^{\lambda z} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K} dz}{\int_0^{\eta-\Delta} e^{\lambda z} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K} dz}. \end{aligned}$$

Suppose $\Delta < 0$

$$\begin{aligned} E[e^{-g(\Delta+z)}|D_\Delta, t_i] &= \frac{\int_{-\Delta}^{\eta} e^{-g(z+\Delta)} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K} e^{\lambda z} dz}{\int_{-\Delta}^{\eta} e^{\lambda z} \frac{(I-1)!}{(K-1)!(I-1-K)!} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K} dz} \\ &= \frac{\int_{-\Delta}^{\eta} e^{-g(z+\Delta)} e^{\lambda z} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K} dz}{\int_{-\Delta}^{\eta} e^{\lambda z} (z+\Delta)^{K-1} (\eta-(z+\Delta))^{I-1-K} dz}. \end{aligned}$$

A.2.3 Simplifying τ

For τ strategies to comprise a symmetric equilibrium, it must be the case that i can do no better than to choose $\Delta = 0$. A necessary condition is that $h(0, t_i) E[1 - e^{-g(z+\tau)}|D_0, t_i] = g$. Solving this expression for τ yields

$$\tau = \frac{1}{g} \ln \left(\frac{E[e^{-g(z)}|D_0, t_i]}{1 - \frac{g}{h(0, t_i)}} \right).$$

Recall that, after cancellation, we can rewrite the above expression as:

$$\frac{E[e^{-g(z)}|D_0, t_i]}{1 - \frac{g}{h(0, t_i)}} = \frac{\int_0^{\eta} e^{-(g-\lambda)z} (z)^{K-1} (\eta-z)^{I-1-K} dz}{\int_0^{\eta} e^{\lambda z} (z)^{K-1} (\eta-z)^{I-1-K} dz - g \int_0^{\eta} e^{\lambda z} \sum_{n=0}^{K-1} \frac{(K-1)!(I-1-K)!}{n!(I-1-n)!} (z)^n (\eta-z)^{I-1-n} dz}. \quad (9)$$

Now, let Num and Den denote, respectively, the numerator and denominator of the right-hand side of equation (9). By series expansion of the exponential function it follows that

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

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

In the expression for Den ,

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

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

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

Hence,

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

Therefore,

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

Next, we show that for $\eta \in [0, \bar{\eta}]$, $\tau > 0$. Since $E[e^{-g(z)} | D_0, t_i] > 0$, the following lemma is sufficient.

Lemma 1 $h(0, t_i) > g$.

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

$$\begin{aligned} \frac{E[e^{-g(z)}|D_0, t_i]}{1 - \frac{g}{h(0, t_i)}} &= \frac{\lambda F(K, I, \eta(\lambda - g))}{g - (g - \lambda) F(K, I, \eta\lambda)} \\ &> \frac{\lambda F(K, I, \eta(\lambda - g))}{g - (g - \lambda) F(K, I, \bar{\eta}\lambda)} \\ &= \frac{\lambda F(K, I, \eta(\lambda - g))}{g - (g - \lambda) \left(\frac{I g}{I g - (I - K + 1)\lambda} \right)} \\ &> F(K, I, \eta(\lambda - g)) > 0. \blacksquare \end{aligned}$$

A.2.4 Global Deviation

We are now in a position to show that Δ strategies do not constitute profitable global deviations. A necessary condition for an equilibrium is that the “marginal cost” for the $\Delta = 0$ strategy is equal to the (constant) “marginal benefit” of g . To show that there is no profitable global deviation, it is sufficient to show that the marginal costs are increasing in Δ .

The marginal cost of a Δ strategy (conditional on t_i) is

$$\begin{aligned} MC(\Delta|t_i) &= h(\Delta, t_i) (1 - e^{-g\tau} E[e^{-g(t_i + \Delta - t_0)}|D_\Delta, t_i]) \\ &= h(\Delta, t_i) \left(1 - \frac{h(0, t_i) - g}{h(0, t_i) E[e^{-g(t_i - t_0)}|D_{t_i}, t_i]} E[e^{-g(t_i + \Delta - t_0)}|D_\Delta, t_i] \right) \\ &= \frac{h(\Delta, t_i)}{h(0, t_i)} \left(h(0, t_i) + (g - h(0, t_i)) \left(\frac{E[e^{-g(t_i + \Delta - t_0)}|D_\Delta, t_i]}{E[e^{-g(t_i - t_0)}|D_{t_i}, t_i]} \right) \right). \end{aligned}$$

It sufficient to show that $h(\Delta, t_i)$ is increasing in Δ and $E[e^{-g(t_i + \Delta + \tau - t_0)}|D_\Delta, t_i]$ is decreasing in Δ .

Claim 1: Suppose that $\Delta \in [0, \eta]$, then $h(\Delta, t_i)$ is increasing in Δ .

Proof: First, define $a := z + \Delta$. Then we can rewrite

$$h(\Delta, t_i) = \frac{\frac{(I-1)!}{(K-1)!(I-1-K)!} \int_{\Delta}^{\eta} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da}{\sum_{n=0}^{K-1} \binom{I-1}{n} \int_{\Delta}^{\eta} e^{\lambda a} (a)^n (\eta - (a))^{I-1-n} da}.$$

Next, differentiate $h(\Delta, t_i)$ with respect to Δ . The sign of this derivative takes on the sign of

$$\begin{aligned}
& -e^{\lambda\Delta} (\Delta)^{K-1} (\eta - \Delta)^{I-1-K} \sum_{n=0}^{K-1} \binom{I-1}{n} \int_{\Delta}^{\eta} e^{\lambda a} (a)^n (\eta - (a))^{I-1-n} da \\
& + \sum_{n=0}^{K-1} \binom{I-1}{n} e^{\lambda\Delta} (\Delta)^n (\eta - \Delta)^{I-1-n} \int_{\Delta}^{\eta} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da.
\end{aligned}$$

We can then rewrite this expression as

$$\begin{aligned}
& - \sum_{n=0}^{K-1} \binom{I-1}{n} (\Delta)^{K-1} (\eta - \Delta)^{I-1-K} \int_{\Delta}^{\eta} e^{\lambda a} (a)^n (\eta - a)^{I-1-n} da \\
& + \sum_{n=0}^{K-1} \binom{I-1}{n} (\Delta)^n (\eta - \Delta)^{I-1-n} \int_{\Delta}^{\eta} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da \\
= & - \sum_{n=0}^{K-1} \binom{I-1}{n} (\Delta)^n (\eta - \Delta)^{I-1-K} \int_{\Delta}^{\eta} e^{\lambda a} (\Delta)^{K-1-n} (a)^n (\eta - a)^{I-1-n} da \\
& + \sum_{n=0}^{K-1} \binom{I-1}{n} (\Delta)^n (\eta - \Delta)^{I-1-K} \int_{\Delta}^{\eta} e^{\lambda a} (a)^{K-1} (\eta - \Delta)^{K-n} (\eta - a)^{I-1-K} da > 0,
\end{aligned}$$

where the inequality follows from the fact that $a > \Delta$ and $\eta - \Delta > \eta - a$ for almost all a . \square

Claim 2: Suppose that $\Delta \in [0, \eta]$, then $E [e^{-g(z+\Delta-t_0)} | D_{\Delta}, t_i]$ is decreasing in Δ .

Proof. Recall that

$$E [e^{-g(z+\Delta-t_0)} | D_{\Delta}, t_i] = \frac{\int_{\Delta}^{\eta} e^{-ga} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da}{\int_{\Delta}^{\eta} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da}$$

Next, differentiate $E [e^{-g(z+\Delta-t_0)} | D_{\Delta}, t_i]$ with respect to Δ . The sign of this derivative takes on the sign of

$$\begin{aligned}
& -e^{-g\Delta} e^{\lambda\Delta} (\Delta)^{K-1} (\eta - \Delta)^{I-1-K} \int_{\Delta}^{\eta} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da \\
& + e^{\lambda\Delta} (\Delta)^{K-1} (\eta - \Delta)^{I-1-K} \int_{\Delta}^{\eta} e^{-ga} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da \\
\propto & \int_{\Delta}^{\eta} -e^{-g\Delta} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da + \int_{\Delta}^{\eta} e^{-ga} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da \\
= & \int_{\Delta}^{\eta} (e^{-ga} - e^{-g\Delta}) e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da < 0,
\end{aligned}$$

where the inequality follows from the fact that $a > \Delta$. \square

Claim 3: Suppose that $\Delta \in [-\max\{\eta, \tau\}, 0]$. Then $h(\Delta, t_i)$ is increasing in Δ .

Proof. First, define $b := \eta - (z + \Delta)$, then we can rewrite:

$$\begin{aligned}
h(\Delta, t_i) &= \frac{\frac{(I-1)!}{(K-1)!(I-1-K)!} \int_{-\Delta}^{\eta} e^{\lambda(\eta-\Delta-b)} (\eta-b)^{K-1} (b)^{I-1-K} db}{\int_{-\Delta}^{\eta} e^{\lambda(\eta-\Delta-b)} \sum_{n=0}^{K-1} \binom{I-1}{n} (\eta-b)^n (b)^{I-1-n} db + \eta^{I-1} \frac{e^{-\lambda\Delta}-1}{\lambda}} \\
&= \frac{\frac{(I-1)!}{(K-1)!(I-1-K)!} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta-b)^{K-1} (b)^{I-1-K} db}{\int_{-\Delta}^{\eta} e^{-\lambda b} \sum_{n=0}^{K-1} \binom{I-1}{n} (\eta-b)^n (b)^{I-1-n} db + e^{-\lambda(\eta-\Delta)} \eta^{I-1} \frac{e^{-\lambda\Delta}-1}{\lambda}} \\
&= \frac{\frac{(I-1)!}{(K-1)!(I-1-K)!} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta-b)^{K-1} (b)^{I-1-K} db}{\sum_{n=0}^{K-1} \binom{I-1}{n} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta-b)^n (b)^{I-1-n} db + \eta^{I-1} e^{-\lambda\eta} \frac{1-e^{\lambda\Delta}}{\lambda}}.
\end{aligned}$$

Next, differentiate $h(\Delta, t_i)$ with respect to Δ . The sign of this derivative takes on the sign of

$$\begin{aligned}
&e^{\lambda\Delta} (\eta + \Delta)^{K-1} (-\Delta)^{I-1-K} \left[\sum_{n=0}^{K-1} \binom{I-1}{n} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta-b)^n (b)^{I-1-n} db + \eta^{I-1} e^{-\lambda\eta} \frac{1-e^{\lambda\Delta}}{\lambda} \right] \\
&- \left[\sum_{n=0}^{K-1} \binom{I-1}{n} e^{\lambda\Delta} (\eta + \Delta)^n (-\Delta)^{I-1-n} - \eta^{I-1} e^{-\lambda\eta} e^{\lambda\Delta} \right] \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta-b)^{K-1} (b)^{I-1-K} db
\end{aligned}$$

First, note that two the non-sum terms are positive.

$$\begin{aligned}
&e^{\lambda\Delta} (\eta + \Delta)^{K-1} (-\Delta)^{I-1-K} \eta^{I-1} e^{-\lambda\eta} \frac{1-e^{\lambda\Delta}}{\lambda} + \eta^{I-1} e^{-\lambda\eta} e^{\lambda\Delta} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta-b)^{K-1} (b)^{I-1-K} db \\
&= e^{\lambda\Delta} e^{-\lambda\eta} \eta^{I-1} \left((\eta + \Delta)^{K-1} (-\Delta)^{I-1-K} \frac{1-e^{\lambda\Delta}}{\lambda} + \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta-b)^{K-1} (b)^{I-1-K} db \right) > 0,
\end{aligned}$$

since $1 - e^{\lambda\Delta} > 0$.

Next, we show that the sum terms are positive.

$$\begin{aligned}
& e^{\lambda\Delta} (\eta + \Delta)^{K-1} (-\Delta)^{I-1-K} \sum_{n=0}^{K-1} \binom{I-1}{n} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta - b)^n (b)^{I-1-n} db \\
& - \sum_{n=0}^{K-1} \binom{I-1}{n} e^{\lambda\Delta} (\eta + \Delta)^n (-\Delta)^{I-1-n} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta - b)^{K-1} (b)^{I-1-K} db \\
= & \sum_{n=0}^{K-1} \binom{I-1}{n} e^{\lambda\Delta} (\eta + \Delta)^{K-1} (-\Delta)^{I-1-K} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta - b)^n (b)^{I-1-n} db \\
& - \sum_{n=0}^{K-1} \binom{I-1}{n} e^{\lambda\Delta} (\eta + \Delta)^n (-\Delta)^{I-1-n} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta - b)^{K-1} (b)^{I-1-K} db \\
= & \sum_{n=0}^{K-1} \binom{I-1}{n} e^{\lambda\Delta} (\eta + \Delta)^n (-\Delta)^{I-1-K} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta + \Delta)^{K-1-n} (\eta - b)^n (b)^{I-1-n} db \\
& - \sum_{n=0}^{K-1} \binom{I-1}{n} e^{\lambda\Delta} (\eta + \Delta)^n (-\Delta)^{I-1-K} \int_{-\Delta}^{\eta} e^{-\lambda b} (\eta - b)^{K-1} (-\Delta)^{K-n} (b)^{I-1-K} db > 0,
\end{aligned}$$

where the inequality follows from the fact that $b > -\Delta$. \square

Claim 4: Suppose that $\Delta \in [-\max\{\eta, \tau\}, 0]$, then $E[e^{-g(z+\Delta-t_0)} | D_{\Delta}, t_i]$ is decreasing in Δ .

Proof. Recall that

$$\begin{aligned}
E[e^{-g(z+\Delta-t_0)} | D_{\Delta}, t_i] &= \frac{\int_0^{\eta+\Delta} e^{-g(a)} e^{\lambda(a-\Delta)} (a)^{K-1} (\eta - a)^{I-1-K} da}{\int_0^{\eta+\Delta} e^{\lambda(a-\Delta)} (a)^{K-1} (\eta - a)^{I-1-K} da} \\
&= \frac{\int_0^{\eta+\Delta} e^{-g(a)} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da}{\int_0^{\eta+\Delta} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da}
\end{aligned}$$

Next, differentiate $E[e^{-g(z+\Delta-t_0)} | D_{\Delta}, t_i]$ with respect to Δ . The sign of this derivative takes on the sign of

$$\begin{aligned}
& e^{-g(\eta+\Delta)} e^{\lambda(\eta+\Delta)} (\eta + \Delta)^{K-1} (-\Delta)^{I-1-K} \int_0^{\eta+\Delta} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da \\
& - e^{\lambda(\eta+\Delta)} (\eta + \Delta)^{K-1} (-\Delta)^{I-1-K} \int_0^{\eta+\Delta} e^{-g(a)} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da \\
\propto & \int_0^{\eta+\Delta} e^{-g(\eta+\Delta)} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da - \int_0^{\eta+\Delta} e^{-g(a)} e^{\lambda a} (a)^{K-1} (\eta - a)^{I-1-K} da < 0,
\end{aligned}$$

where the inequality follows from the fact that $a < \eta + \Delta$. \square

Suppose that $\Delta > \eta$. In that case, i exits after $t_i + \eta + \tau$, at which point the game has already ended with certainty. This strategy is dominated by exiting at t_i . Finally, suppose i chooses to exit before becoming informed. Notice that $h(-\tau, t_i) > \underline{h}$ where \underline{h} is the hazard rate associated with exiting before getting a signal; hence, exiting at t_i dominates exiting before receiving a signal.

A.2.5 Uniqueness of Symmetric Equilibria

So far, we have shown the τ strategies comprise a symmetric equilibrium in clock games. Next, we establish that these strategies constitute the unique symmetric equilibrium. Suppose not, then there exists some other delay time $\tau' \neq \tau$ such that the marginal benefit=marginal cost condition is satisfied. That is, given that all players wait τ' periods,

$$h_{t_i+\tau'} E \left[1 - e^{-g(t_i+\tau'-t_0)} \middle| \cdot \right] = g.$$

where $h_{t_i+\tau'}$ represents the hazard rate when all players including i play τ' strategies. Notice that since equilibrium play involves merely symmetrically shifting the delay time from τ to τ' periods, the hazard rate is unchanged. That is,

$$h_{t_i+\tau'} = h(0, t_i).$$

The term $E \left[1 - e^{-g(t_i+\tau'-t_0)} \middle| \cdot \right]$; however, is increasing in τ' ; therefore, for all $\tau' \neq \tau$, it immediately follows that

$$h_{t_i+\tau'} E \left[1 - e^{-g(t_i+\tau'-t_0)} \middle| \cdot \right] = h(0, t_i) E \left[1 - e^{-g(t_i+\tau'-t_0)} \middle| \cdot \right] \neq g.$$

This completes the proof of Proposition 1. \blacksquare

A.3 Proof of Proposition 2

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

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

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

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

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

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

A.4 Proof of Proposition 4

It is straightforward to obtain expressions for $h_1(0, t_i)$ and $E[e^{-g(z)} | D_0, t_i]$. Simply use the analogous expressions given in the proof of Proposition 1 and substitute $K = 1$. This yields

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

Using steps analogous to the simplification of τ , we have

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

Since the monotonicity conditions in the proof of Proposition 1 continue to hold in the present setting, global deviations from τ_1 strategies are not profitable. It remains only to show that $\tau_1 > 0$. Since $E[e^{-g(z)} | D_0, t_i] > 0$, the required inequality follows from

Lemma 2 $h_1 > \frac{Ig}{I-K+1}$.

Proof. Following the identical steps in the proof of Lemma 1, we obtain

$$\frac{E[e^{-g(z)} | D_0, t_i]}{1 - \frac{g_1}{h_1}} = \frac{\lambda F(1, I, \eta(-g+\lambda))}{\left(\lambda - \frac{Ig}{I-K+1} \right) F(1, I, \eta\lambda) + \frac{Ig}{I-K+1}} \geq 0, \quad (10)$$

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. ■

A.5 Proof of Proposition 3

Proof. Following observing the first exit, the identity of the last “type” to exit in the continuation game is commonly known. Formally, fix a perfect Bayesian equilibrium in the continuation game following the first exit. Let $h(t_i)$ be the (endogenous) hazard rate of the game ending associated with type t_i . Define the set of “last” types as:

$$L = \{t_i : h(t_i) > k\},$$

where k is some sufficiently large but finite constant amount. Notice that there is at least one commonly known type facing an infinite hazard rate. Moreover, there is a positive measure of types facing arbitrarily large hazard rates. These types are in the set L . Notice that, since the (growth) benefit of waiting an additional period is at most g , these types have a profitable deviation of exiting before the time indicated in the equilibrium. But this is a contradiction. Therefore, the unique continuation PBE is where all types exit immediately after observing the first exit.

A.6 Proofs of Subsection 6

Lemma 3 $\tau(K - 1, I) < \tau(K, I)$.

Proof. The inequality follows from the fact that $F(a, b, x)$ is increasing in its first argument. ■

Proof of Proposition 5

From Lemma 3, the interval $[\tau(K - 1, I), \tau(K, I)]$ is non-empty. To see that $\tau_{CC} \in [\tau(K - 1, I), \tau(K, I)]$ comprise a symmetric equilibrium, fix the strategies of all players other than i at $\tau_{CC} \in [\tau(K - 1, I), \tau(K, I)]$ and consider a deviation by player i to some time $t_i + \tau_{CC} + \Delta$. Using reasoning identical to that used in the proof of Proposition 1, one can show that if local deviations are not profitable then global deviations are likewise not profitable. Thus, it suffices to show that local deviations are not profitable. This requires

$$\lim_{\Delta \nearrow 0} h(\Delta, t_i) (1 - e^{-g\tau_{CC}} E[e^{-g(t_i + \Delta - t_0)} | D_{\Delta}, t_i]) \leq g$$

and

$$\lim_{\Delta \searrow 0} h(\Delta, t_i) (1 - e^{-g\tau_{CC}} E[e^{-g(t_i + \Delta - t_0)} | D_{\Delta}, t_i]) \geq g.$$

The reason for taking left and right hand limits separately is that a positive mass of players of size $1/I$ are known by player i to be exiting at $t_i + \tau_{CC}$.

First consider a deviation from below.

$$\begin{aligned}
& \lim_{\Delta \nearrow 0} h(\Delta, t_i) (1 - e^{-g\tau_{CC}} E [e^{-g(t_i + \Delta - t_0)} | D_\Delta, t_i]) \\
&= \lambda \frac{F(K, I, \eta\lambda) - e^{-g\tau_{CC}} F(K, I, (\lambda - g)\eta)}{[F(K, I, \eta\lambda) - 1]} \\
&\leq \lambda \frac{F(K, I, \eta\lambda) - e^{-g\tau(K, I)} F(K, I, (\lambda - g)\eta)}{[F(K, I, \eta\lambda) - 1]} = g,
\end{aligned}$$

where the inequality follows from the definition of $\tau(K - 1, I)$ and the fact that $\lim_{\Delta \nearrow 0} h(\Delta, t_i) (1 - e^{-g\tau_{CC}} E [e^{-g(t_i + \Delta - t_0)} | D_\Delta, t_i])$ is increasing in τ_{CC} .

Next, notice that

$$\begin{aligned}
& \lim_{\Delta \searrow 0} h(\Delta, t_i) (1 - e^{-g\tau_{CC}} E [e^{-g(t_i + \Delta - t_0)} | D_\Delta, t_i]) \\
&= \lambda \frac{F(K - 1, I, \eta\lambda) - e^{-g\tau_{CC}} F(K - 1, I, (\lambda - g)\eta)}{[F(K - 1, I, \eta\lambda) - 1]} \\
&\geq \lambda \frac{F(K - 1, I, \eta\lambda) - e^{-g\tau(K-1, I)} F(K - 1, I, (\lambda - g)\eta)}{[F(K - 1, I, \eta\lambda) - 1]} = g,
\end{aligned}$$

where the inequality follows from the definition of $\tau(K, I)$ and the fact that $\lim_{\Delta \searrow 0} h(\Delta, t_i) (1 - e^{-g\tau_{CC}} E [e^{-g(t_i + \Delta - t_0)} | D_\Delta, t_i])$ is increasing in τ_{CC} . Therefore, τ_{CC} strategies comprise a symmetric equilibrium. ■

Proof of Proposition 6

Define $\kappa \equiv \frac{K}{I}$. It is sufficient to show

$$\frac{\lambda F(\kappa I, I, -\eta(g - \lambda))}{g - (g - \lambda) F(\kappa I, I, \eta\lambda)} > \frac{\lambda e^{-\kappa\eta(g - \lambda)}}{g - (g - \lambda) e^{\kappa\eta\lambda}}.$$

To obtain the required inequality, it suffices to show $F(\kappa I, I, x) > e^{\kappa x}$ for $x \neq 0$ and $\kappa \in (0, 1)$. When $x > 0$, $F(\kappa I, I, x)$

$$\begin{aligned}
&= 1 + \frac{\kappa I}{I} x + \frac{\kappa I (\kappa I + 1) x^2}{I(I + 1) 2!} + \frac{\kappa I (\kappa I + 1) (\kappa I + 2) x^3}{I(I + 1) (I + 2) 3!} + \dots \\
&= 1 + \frac{\kappa I}{I} x + \frac{\kappa I (\kappa (I + 1) + (1 - \kappa)) x^2}{I(I + 1) 2!} + \frac{\kappa I (\kappa (I + 1) + (1 - \kappa)) (\kappa (I + 2) + 2(1 - \kappa)) x^3}{I(I + 1) (I + 2) 3!} + \dots \\
&> 1 + \frac{\kappa I}{I} x + \frac{\kappa I (\kappa (I + 1)) x^2}{I(I + 1) 2!} + \frac{\kappa I \kappa (I + 1) (\kappa (I + 2)) x^3}{I(I + 1) (I + 2) 3!} + \dots = e^{\kappa x}.
\end{aligned}$$

For $x < 0$,

$$F(\kappa I, I, x) = e^x F((1 - \kappa) I, I, -x).$$

Using the same steps as above, note $F((1 - \kappa) I, I, -x) > e^{-(1 - \kappa)x}$. Hence, $F(\kappa I, I, x) > e^{\kappa x}$. ■

Proof of Proposition 7

Recall that,

$$\lim_{I \rightarrow \infty} F(\kappa I, I, x) = 1 + \kappa x + \kappa^2 \frac{x^2}{2!} + \kappa^3 \frac{x^3}{3!} + \dots = e^{\kappa x}.$$

Hence,

$$\begin{aligned} \lim_{I \rightarrow \infty} \tau(\kappa I, I) &= \lim_{I \rightarrow \infty} \left(\frac{1}{g} \ln \frac{F(\kappa I, I, \eta(\lambda - g))}{\frac{g}{\lambda} - \frac{g-\lambda}{\lambda} F(\kappa I, I, \eta\lambda)} \right), \\ &= \frac{1}{g} \ln \left(\frac{\lambda e^{\kappa\eta(\lambda-g)}}{g - (g-\lambda) e^{\kappa\eta\lambda}} \right) = \tau_{AB}. \end{aligned}$$

Similarly, $\lim_{I \rightarrow \infty} F(\kappa I - 1, I, x) = e^{\kappa x}$ and hence,

$$\lim_{I \rightarrow \infty} \tau(\kappa I - 1, I) = \tau_{AB} \blacksquare$$

Proof of Proposition 8

Recall that τ_1 is an equilibrium when there is a discrete number, I , of players. From the series expansion of F , we have that $F(0, I, x) = 1$. Then, using steps identical to those in Proposition 5 yields the result. \blacksquare

Proof of Proposition 9

First notice that Proposition 3 directly applies; therefore, all remaining players immediately exit following the first player's exit. Next, fix an equilibrium where $\tau_{1,CC} \geq 0$. Notice that, in such an equilibrium, the hazard rate, h_1 , at time $t_i + \tau_{1,CC}$ must be infinite. This is because another player is certain to exit at the next instant. Further, at time $t_i + \tau_{1,CC}$, player i also infers that he is the first person to receive the signal (otherwise the game would already have ended) hence $t_i = t_0$. Together, these two observations imply that

$$\tau_{1,AB} = \frac{1}{g} \left[\ln \frac{h_1}{h_1 - g \frac{I-1}{I-K}} + \ln E \left[e^{-g(t_i - t_0)} | D_0, t_i \right] \right] = 0 \blacksquare$$

Proof of Proposition 10. Define $\kappa \equiv \frac{K}{I}$ and recall that $\lim_{b \rightarrow \infty} F(1, b, x) = 1$

$$\begin{aligned} \lim_{I \rightarrow \infty} \tau_1 &= \lim_{I \rightarrow \infty} \frac{1}{g} \ln \left(\frac{\lambda F(1, I, \eta(-g+\lambda))}{\frac{Ig}{I(1-\kappa)+1} - \left(\frac{Ig}{I(1-\kappa)+1} - \lambda \right) F(1, I, \eta\lambda)} \right) \\ &= \frac{1}{g} \ln(1) = 0 \blacksquare \end{aligned}$$

Appendix B: Instructions

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

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 B1). 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 B2).

[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 B3).

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 half a 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.

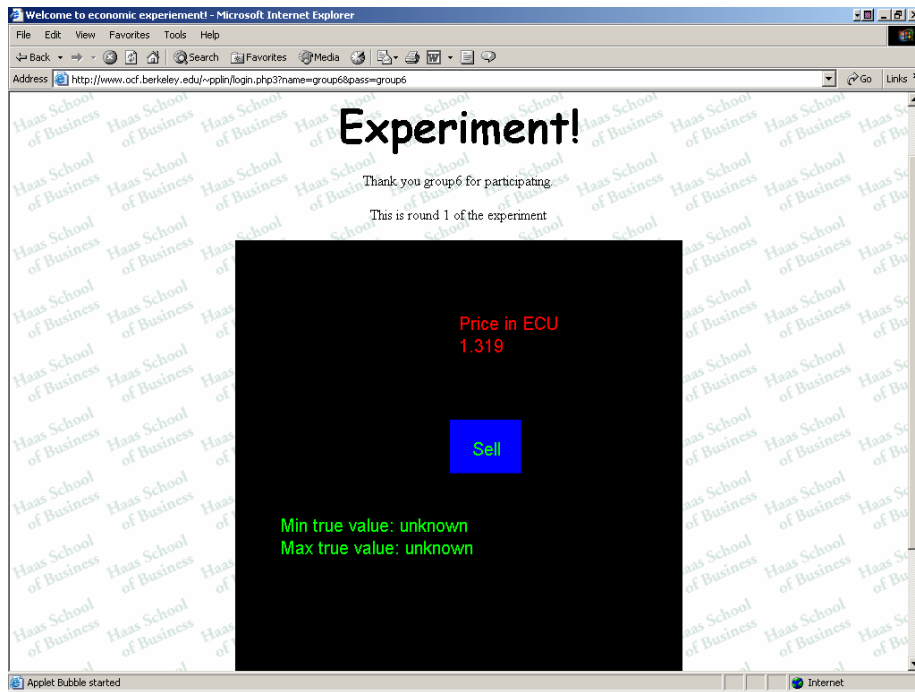


Figure B1




Figure B2

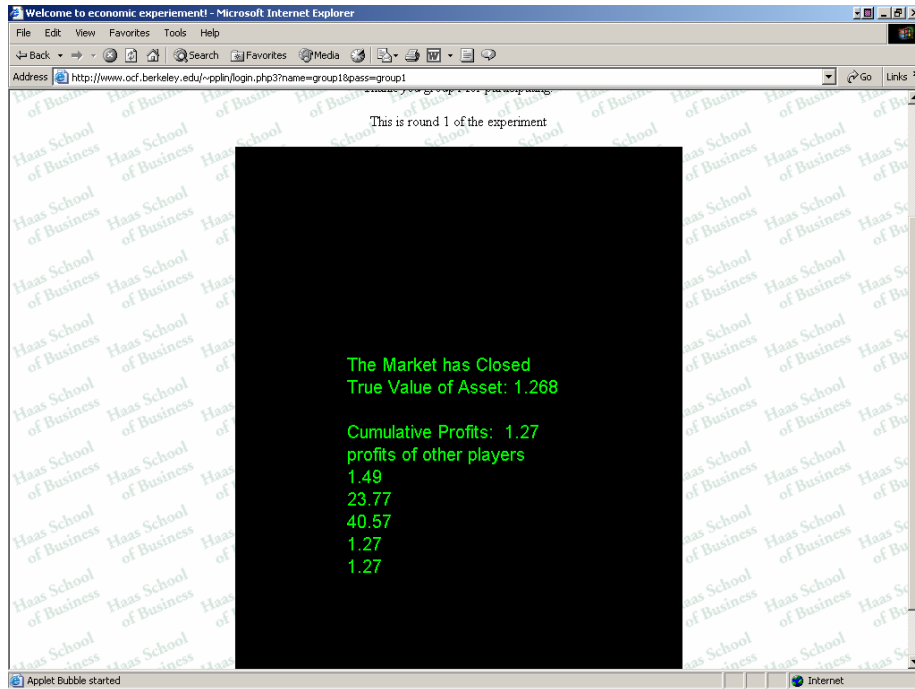


Figure B3