Characterizing Persistent Disequilibrium Dynamics: Imitation or Optimization?*

Daniel G. Stephenson† Department of Economics Texas A&M University

Alexander L. Brown‡ Department of Economics Texas A&M University

May 24, 2016

Abstract

A large class of adaptive models take imitation as the primary diver of behavior, while others focus on some form of myopic optimization. Previous literature is mixed on which class of adaptive models is more appropriate for describing human behavior. We conduct an experiment using a continuous-time, all-pay auction, providing richer data than previous studies and clean separation between the theoretical predictions of imitative models and those of optimization models. In accordance with theoretical predictions from adaptive models, but in contrast to Nash equilibrium predictions, we observe incessant disequilibrium cycles in bidding behavior. Myopic optimization models greatly outperform imitative models in characterizing the observed behavior. The provision of social information regarding the bids and payoffs of others—a required informational assumption for imitative models—neither increases the predictive power of imitative models nor disrupts the observed bidding cycles. Instead, it increases both the speed of cycles and the precision of optimization behavior.

*We thank the Texas A&M Humanities and Social Science Enhancement of Research Capacity Program and the College of Liberal Arts at Texas A&M for their generous financial support of this research. Furthermore, we thank Catherine C. Eckel, and the experimental research team at Texas A&M for their continual support and feedback. We are also grateful for the insightful comments from attendees of the 2013 Economic Science Association World Meetings.

†Corresponding author. Please contact at dgstephenson@tamu.edu.

‡Email: alexbrown@tamu.edu.

1 Introduction

Nash Equilibrium is a powerful tool for understanding strategic behavior. Even when agents fall short of perfect rationality, simple adaptive processes can often drive long-run behavior towards equilibrium predictions. For this reason, adaptive models have long been employed to justify the application of equilibrium solution concepts in the presence of bounded rationality (e.g., Cournot et al., 1897; Nash, 1951). In such environments, the long run predictions of many adaptive models are remarkably similar to equilibrium, making it difficult to separate the predicitons of alternative adaptive models. However, in strategic environments with unstable equilibria, adaptive dynamics may fail to converge, leading boundedly rational agents to exhibit persistent disequilibrium behavior. In these unstable environments, we can clearly separate the distinct theoretical predictions from different adaptive models, as behavior remains strongly non-convergent. Moreover, they provide strong evidence that characterizing the adaptive processes from which equilibria emerge can be equally important as characterizing equilibrium behavior.[1]

Different adaptive models often make fundamentally different behavioral assumptions. In particular, a large class of adaptive models take imitation as the primary diver of adaptive behavior, while others focus on some form of myopic optimization. To investigate the explanatory power of these distinct classes of adaptive models, this paper presents experimental results from laboratory experiments on continuous-time all-pay auctions[2] with continuous strategy spaces, discontinuous payoff functions, a finite population of agents, and highly a unstable Nash equilibrium. These combined features foster disequilibrium bidding cycles, an outcome running contrary to the predictions of both the Nash and quantal response equilibrium but consistent with adaptive models. Further, the cycles, themselves, generate

[1]For more on the importance of understanding adaptive dynamics out of equilibrium, see Sandholm (2010) or Weibull (1997).

[2]The all-pay auction has been examined extensively in experimental environments (see Dechenaux et al., 2014, for a survey). Previous experimental studies (e.g., Gneezy and Smorodinsky, 2006; Lugovskyy et al., 2010; Ernst and Thöni, 2013) of the all-pay auction conduct a sequence of discrete rounds in which subjects secretly select their bids and the single highest bidder receives a price. Unlike these previous studies, we investigate all-pay auctions with three bidders who compete over two equally valuable prizes. The top two bidders each receive a prize, while the lowest bidder receives no prize. As in the standard all-pay auction, every bidder pays her bid. In contrast to previous studies, our subjects adjust their bids asynchronously throughout the experimental session and earn flow payoffs continuously over time.

distinctly different, instantaneous predictions of imitative and myopic optimization models.

Imitative models provide a convenient explanation for the predictive power of equilibrium without resorting to the assumption of perfectly rational agents. If agents imitate the most successful strategies, then other strategies will gradually die out, and successful strategies will be increasingly employed by the population. Such imitative dynamics illustrate how long run behavior can stabilize on equilibrium even when agents have little understanding of the strategic incentives they face. For example, Vega-Redondo (1997) proves that a noisy imitate-the-best dynamic leads to the emergence of Walrasian equilibrium under Cournot competition between a finite number of firms. However, this explanation of equilibrium remains unorthodox in in economics, where agents are typically assumed to be goal-oriented utility maximizers. Accordingly, other models of adaptive behavior describe boundedly rational agents who myopically respond to their individual payoff incentives instead of imitating the successful strategies employed by others (e.g., Fudenberg and Levine, 1998; Gilboa and Matsui, 1991; Smith, 1984).

Each of these adaptive models place different informational requirements on agents. Imitative dynamics describe agents who mimic the actions taken by their successful peers, so imitation necessarily requires that agents observe the actions and payoffs of others. In contrast, myopic optimization dynamics describe agents who select actions that maximize their individual payoffs, so the minimum information requirement for agents under these models is simply to observe their own payoffs. To examine the robustness of adaptive models to these informational requirements, our treatments manipulate the availability of social information regarding the bids and payoffs of others. In our local information treatment, subjects can only observe their own bids and payoffs. In our global information treatment, subjects can also observe the bids and payoffs of others. We hypothesize that imitative models will model subject behavior well when social information is present, while myopic optimization models will work in the absence of social information.

Contrary to this hypothesis, when social information is provided for subjects, optimization models provide more reliable predictions of bidding behavior than imitative models. Specifically, we find that optimization behavior is more precise, bids are higher, earnings are lower, and bidding cycles occur more rapidly when subjects have access to global information. Subjects with access to social information consistently employ strategies that better respond to their opponents' strategies instead of merely mimicking the most successful strategies employed by others. Under local information, subjects are unable to observe the behavior of others, so they employ a trial-and-error strategy and exhibit significantly more behavioral noise.

The failure of imitative models to reliably characterize subject behavior in this experiment may suggest a general failing of imitative models to accurately represent human cognition. While a direct interpretation of imitative models would confine subjects to playing only previously employed strategies, we consider a more liberal model which predicts that agents are *more likely* to play strategies *close to* the most successful previously employed strategies.[3] Our global-information treatment is designed to make executing this dynamic as easy as possible; subjects only need to click on the highest bar on a computer screen to follow the dynamic exactly (see Section 3.1 for a more detailed description). However, instead of merely imitating successful strategies employed by others, subjects clearly employ more sophisticated optimization methods that involve estimating the forgone payoff to previously unemployed strategies. Moreover, the explanatory power of imitative models changes little with the provision of social information, suggesting that social information helps subjects to more effectively optimize instead of promoting imitation.

Previous literature includes several attempts to investigate whether imitation or optimization best describes actual behavior in games. None have utilized a dataset as rich as ours—our use of continuous time, continuous strategy spaces, and discontinuous payoff functions allows us to reach uniquely definitive conclusions regarding the predictive validity of alternative adaptive dynamics (see section 4 for our specific hypotheses). Perhaps in part because of this data limitation, there appears to be little consensus reached on the degree to which individuals imitate or optimize. Tests of imitative models using a wide variety of experimental designs have have found evidence both in favor (Offerman et al., 2002; Feri et al., 2011) and against (Cheung and Friedman, 1998; Friedman et al., 2015) the predictive power of imitative models. Similarly, research on dynamic cognition and behavioral learning, which like our paper examines how individuals respond to information about games, finds

[3]The deterministic imitate-the-best model is the limiting case of this model.

both support (Offerman et al., 2002) and opposition (Ho et al., 2007; Camerer and Hua Ho, 1999) to the notion that individuals respond to payoff information by imitating previous actions that earned higher payoffs. Our paper's novel experimental design provides conclusive evidence on this issue, strongly rejecting the idea that people are naturally imitators rather than optimizers.

The paper also contributes to the small, but burgeoning area of literature that studies the properties of disequilibrium dynamics in continuous-time games (i.e., Oprea et al., 2011; Cason et al., 2014). Consistent with that literature (specifically, Cason et al.), we observe persistent, disequilibrium cycles that cannot be predicted by standard equilibrium-based models of game theory. Our treatment variable allows us to characterize the channels through which these equilibrium cycles respond to the availability of social information. Specifically, social information does not alleviate disequilibrium cycling, rather it reduces behavioral noise in optimization and increases the frequency of cycles. Further, because our design features a continuous strategy space, it provides detailed data regarding the accuracy with which alternate evolutionary models predict the behavioral dynamics of these disequilibrium cycles. To our knowledge, these contributions are unique within the literature.

This paper proceeds as follows: Section 2 presents the structure of the game and two different equilibrium models. It also describes and the various adaptive models that will be used to characterize the experimental data. Section 3 presents the full design and procedures of the experiment. Section 4 provides our hypotheses. Section 5 presents the main results and Section 6 concludes.

2 Theory

In all-pay auctions, multiple agents expend costly effort to compete over a limited number of prizes. Prizes are awarded to the agents who expend the most effort, but every agent bears the cost of her own effort, even if she does not win a prize. All-pay auctions often model strategic environments that involve both conflict and non-recoverable costs such as political lobbying (Baye et al., 1993), patent races (Marinucci and Vergote, 2011), biological competition (Chatterjee et al., 2012), and international warfare (Hodler and Yektaş, 2012).

The all-pay auction involves three players who compete over two prizes. Each player i starts with a endowment w and selects her bid b_i from the closed interval $[0, w]$. The top two bidders each receive a prize with value v and the lowest bidder receives no prize. Every player must pay her bid, regardless of whether or not she won a prize. In the case of a tie, the winner is determined randomly. Accordingly, the payoff function for player i is given by:

$$\pi_i(b_i, b_j, b_k) = \begin{cases} w - b_i + v & \text{if } b_i > \min\{b_j, b_k\} \\ w - b_i + 2v/3 & \text{if } b_i = b_j = b_k \\ w - b_i + v/2 & \text{if } b_i = \min\{b_j, b_k\} < \max\{b_j, b_k\} \\ w - b_i & \text{otherwise} \end{cases} \quad (1)$$

2.1 Equilibrium Models

We consider equilibrium models including the Nash equilibrium and the logit quantal response equilibrium. Nash equilibrium assumes that each agent selects a best response to the strategies selected by others. In contrast, the logit quantal response equilibrium assumes that agents make probabilistic errors in their payoff evaluations.

2.1.1 Nash Equilibrium

The all-pay auction investigated here with three bidders and two winners has no pure strategy Nash equilibrium, but it does have a unique symmetric mixed strategy Nash equilibrium. First derived by Clark and Riis (1998),[4] the corresponding probability density function for the bid of player i given by

$$f(b_i) = \frac{1}{2v}\left(1 - \frac{b_i}{v}\right)^{-1/2} \qquad \text{for all } b_i \in [0, v]. \quad (2)$$

The black line in Figure 1 illustrates this equilibrium density function. Note that that Nash equilibrium probability density function approaches infinity as bids approaches the value of the prize and remains at zero for any bid above the value of the prize. Thus, in equilibrium players are likely to bid near the value of the prize but never above it.

[4] Appendix Section A.1 contains an alternate derivation of the equilibrium.

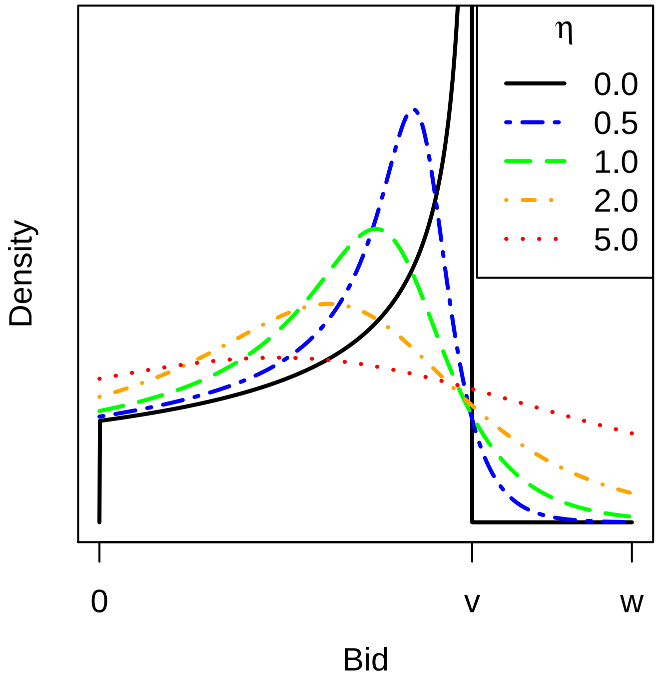

Figure 1: Nash Equilibrium and Logit Quantal Response Equilibria

2.1.2 Logit Quantal Response Equilibrium

While the Nash equilibrium describes the behavior of perfectly rational and perfectly precise agents, the logit quantal response equilibrium described by McKelvey and Palfrey (1995) and Lopez (1995) allows us to model the behavior of imprecise boundedly rational agents. Such agents make probabilistic errors in their evaluation of alternate strategies. Although they typically fail to select a best response, they are more likely to select strategies that yield higher payoffs.

Unlike the perfectly rational agents described by Nash equilibrium, agents in logit quantal response equilibrium may place positive probability on dominated strategies, since their behavior is fundamentally stochastic. In the case of a continuous strategy space, the probability density function for the logit quantal response equilibrium mixed strategy σ_i satisfies:

$$f\left(b_i\right) = \frac{\exp\left(\eta^{-1}\pi_i\left(b_i, \sigma_{-i}\right)\right)}{\displaystyle\int \exp\left(\eta^{-1}\pi_i\left(x, \sigma_{-i}\right)\right) dx} \tag{3}$$

Here η denotes the level of behavioral noise in an agent's evaluation of payoffs. When η is small, agents make small errors, and the strategy distribution approaches the Nash equilibrium. When η is large, the logit quantal response equilibrium approaches uniformly random play. To illustrate this tendency, Figure 1 depicts the logit quantal response equilibrium under alternate values of η.

A closed form solution for the logit quantal response equilibrium of an all-pay auction with a single prize is provided by Anderson et al. (1998). To the best of our knowledge, no closed form solution is currently available for the logit quantal response equilibrium of an all-pay auction with two prizes. Accordingly, a formal derivation of the logit quantal response equilibrium in this case is found in Appendix Section A.2. The logit quantal response equilibrium probability density function for the bid of player i is

$$f\left(b_i\right) = \frac{\eta G\left(\sqrt{\eta v}\right)\exp\left(-\eta b_i\right)}{\sqrt{\eta v}\left[1-\exp\left(-\eta w\right)\right]}\left[\exp\left(G^{-1}\left(G\left(\sqrt{\eta v}\right)\left[1-\frac{1-\exp\left(-\eta b_i\right)}{1-\exp\left(-\eta w\right)}\right]\right)^2\right)\right]^{-1}, \quad (4)$$

where $G\left(x\right) = \int_0^x \exp\left(u^2\right)du = \frac{\sqrt{\pi}}{2}\operatorname{erfi}\left(x\right)$.

2.2 Evolutionary Game Theory

The experiment in this paper involves continuous-time, two good, three-bidder all pay auctions. These auctions are conducted in groups of three, but subjects' rewards are calculated using mean matching, so essentially every subject plays every other subject all the time. In situations like these, it is useful to consider models from evolutionary game theory, the study of "large populations of agents who repeatedly engage in strategic interactions," (Sandholm, 2010).

2.2.1 Evolutionary Stability

By design, this all-pay auction has evolutionary dynamics that make it a prime candidate for persistent disequilibrium. Accordingly, it has no evolutionary stable strategy. The idea of an evolutionarily stable strategy was introduced by Smith and Price (1973), who employed it to identify the stability of biological phenotypes in large populations under the pressures of mutation and natural selection. More recently, game theorists and social scientists have employed evolutionary stability criteria to model the behavioral stability of Nash equilibria in a wide variety of strategic settings.[5]

A strategy is evolutionarily stable if it induces a self-enforcing convention. That is, a strategy x is evolutionarily stable if no other strategy y can invade it when the entire population initially employs strategy x. More formally, in a in a symmetric normal form game, a strategy x is evolutionarily stable if there exists some $C \in (0,1)$ such that for all $\varepsilon \in (0,C)$ and for any other strategy y

$$\pi\,(\,x \mid \epsilon y + (1-\epsilon)\,x\,) > \pi\,(\,y \mid \epsilon y + (1-\epsilon)\,x\,) \tag{5}$$

Thus, if x is evolutionarily stable and a sufficiently small proportion of the population deviates to an alternate strategy y, then agents who employ x will earn a strictly higher payoff than agents who employ y.

The Nash equilibrium strategy for the all-pay auction is not evolutionarily stable. Too see why, suppose that a small proportion ϵ of the population deviates from the Nash equilibrium strategy x to an alternate strategy y under which agents always bid the full value of the prize. Since the support of the equilibrium bid distribution is given by the closed interval $[0,v]$, agents who employ the invading strategy y will win the prize with probability one whenever they are matched against an agent who employs the equilibrium bidding strategy. In this case, the invading strategy y earns a higher expected payoff than the equilibrium mixed strategy x, so the equilibrium mixed strategy for the all-pay auction with three bidders and two prizes is not evolutionarily stable. A formal derivation of this result is found in Appendix Section A.3. Since mixed strategy Nash equilibrium fails to induce a self enforcing convention in this all-pay auction, we expect to observe dynamic instability in experimental bidding behavior.

[5]These settings include price competition (Alos-Ferrer and Ania, 2005), linguistics (Demichelis and Weibull, 2008), and corporate investment (Parayre and Hurry, 2001).

2.2.2 Adaptive Dynamics

As we expect this experimental environment to be rife with instability, we have ample opportunity to examine the adaptive behavior of subjects. In particular, we will specifically examine noisy optimization dynamics and a noisy imitation dynamics from evolutionary game theory (Sandholm, 2010). In these adaptive models, agents make asynchronous strategy adjustments over time. The timing of these adjustments follows a homogeneous Poisson process with a rate of δ adjustments per second. The value b_{it} here denotes the bid employed by agent i at time t. To determine the relative strengths of these models, we also develop a multi-parameter model that nests each as a special case.

Under deterministic optimization models, such as those described by Gilboa and Matsui (1991) and Golman (2011) agents switch precisely to their best response. In contrast, the logit dynamic is a noisy optimization model (Fudenberg and Levine, 1998; Hopkins, 1999), predicting that agents will be more likely to select bids that yield higher payoff. Under this model, the likelihood that an agent i who adjusts her bid at time t will select a particular bid b is given by:

$$f_{i,t}(b) = \frac{\exp\left(\beta\pi_i\left(b, b_{-i,t}\right)\right)}{\int_0^w \exp\left(\beta\pi_i\left(x, b_{-i,t}\right)\right) dx} \tag{6}$$

Purely imitative models (e.g., Duersch et al., 2012; Golman, 2011; Taylor and Jonker, 1978; Oechssler and Riedel, 2001) predict that agents will exclusively imitate the strategies of other agents they encounter. Such models predict that agents will never innovate by playing a strategy that was not previously employed by others in the population. In an experiment such as this one, with a continuous strategy space and finite number of subjects, this prediction will almost certainly fail as subjects pick new strategies that were not previously employed by others. To increase the flexibility of these imitative models, we consider a noisy imitation model under which agents are *more likely* to select bids that are *close* to the bid that is currently employed by the agent with the highest earnings rate. Let b_t^H denote the bid employed by the agent who has the highest earning rate at time t. Under this noisy imitative model, the likelihood that an agent who adjusts her bid at time t will select a particular bid b is given by:

$$f_{i,t}(b) = \frac{\exp\left(\gamma\left|b - b_t^H\right|\right)}{\int_0^w \exp\left(\gamma\left|x - b_t^H\right|\right) dx}. \tag{7}$$

It is important to note that the noiseless imitate-the-best dynamic is a special case of this model where $\gamma \to \infty$.

To examine the relative strength of the noisy imitation and optimization dynamics, we develop a combined model that includes each both imitation and optimization as a special case. In this combined model, the attraction of a bid x for an agent i at time t is given by

$$A_{it}(b) = \alpha \pi_i(b, b_{-i,t}) - \beta |b - b_{i,t}| - \gamma \left| b - b_t^H \right|. \tag{8}$$

Here b_{it} denotes the the bid employed by agent i at time t and b_t^H denotes the bid employed by the agent who is earning the highest payoff at time t. The parameter α denotes the extent to which agents are more likely to select strategies that yield higher payoffs. The parameter β denotes the degree to which bids are autocorrelated, that is, the extent to which agents tend to select bids that are close to their previous bids.[6] The parameter γ captures denotes the tendency to imitate success, that is, the extent to which agents tend to pick bids which are close to the bid employed by the highest earning player. Accordingly, the the likelihood that agent i will select a bid x when she makes an adjustment at time t is given by

$$f_{it}(b) = \frac{\exp(A_{it}(b))}{\int_0^w \exp(A_{it}(x))\,dx}. \tag{9}$$

3 Experimental Design and Procedures

3.1 Design

To implement the game discussed in Section 2, subjects were endowed with $w = \$10$ and competed for prizes with value $v = \$7$. Subject bids were bounded on the interval $[0, w]$. As this game takes place in continuous time, each session consisted of one continuous 40 minute period. During this period, subjects could adjust their bids as frequently as desired with the click of the mouse. Whenever a subject clicked, her bid instantaneously changed to the level corresponding to the horizontal position of her mouse, and the corresponding payoff rates were immediately recalculated.

The experiment consisted of two informational treatments. Under the global information

[6] Results show the removal of this autocorrelation term does not affect the relative explanatory power of

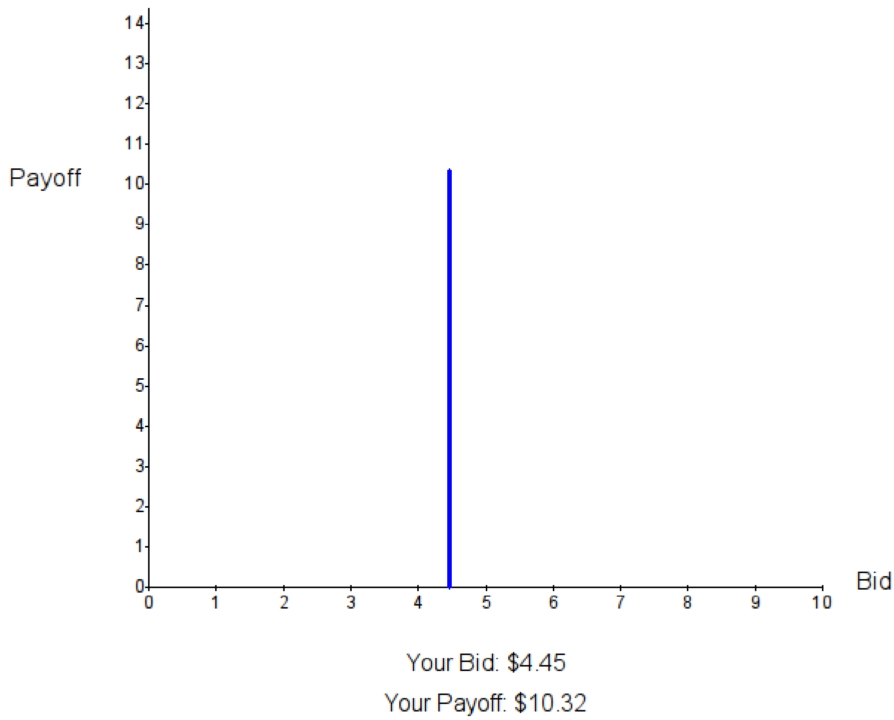

Figure 2: User Interface Under Local Information

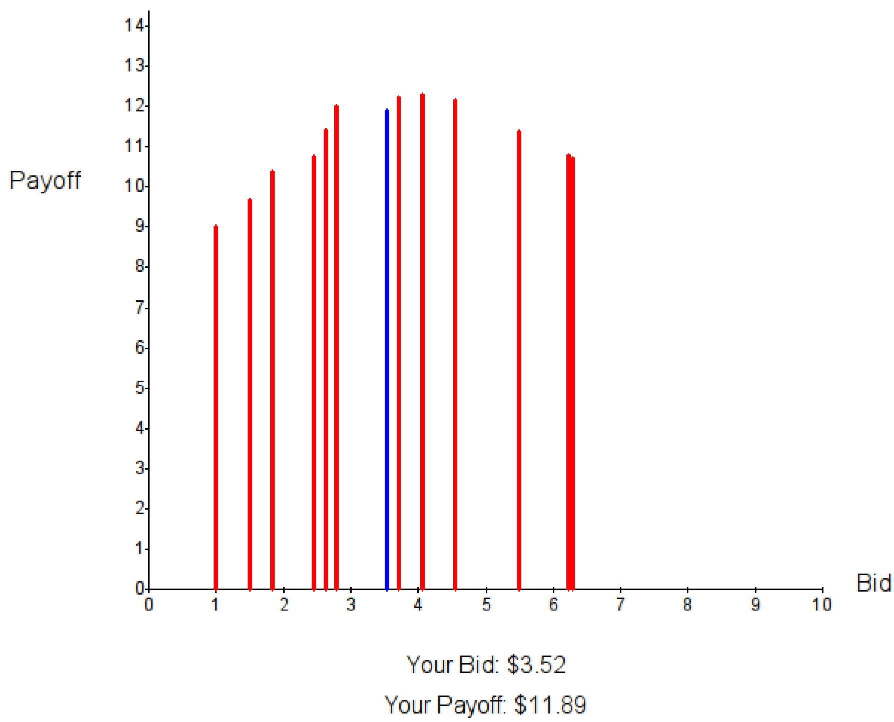

Figure 3: User Interface Under Global Information

treatment, each subject received real-time information regarding the bids and payoffs of every participant in her cohort. Under the local information treatment, subjects only observed their own bid and payoff. In both treatments, bids and payoffs were recorded at a rate of four times per second.

Figures 2 and 3 illustrate the experimental interface under the local-information and global-information treatments, respectively. The subject's current bid and payoff is represented by a blue line. The horizontal position of the blue line indicates the subject's current bid and the height of the blue line indicates the subject's current payoff. The subject's current bid and payoff are displayed numerically at the bottom of the screen. Under the global-information treatment, the current bid and payoff of each other subject is represented by a red line.

To provide random rematching in continuous time, we employ a mean matching protocol (e.g. Cason et al., 2014; Oprea et al., 2011). Each subject's instantaneous payoff is given by the expected value of her payoff from being randomly matched into a group of three agents. By the law of large numbers, high frequency mean-matching provides a superior approximation to truly continuous random matching than does high frequency random matching.

3.2 Procedures

Thirty subjects participated in one session of the global-information treatment and 27 subjects participated in one session of the local-information treatment. Subjects were recruited from the Texas A&M undergraduate population using `econdollars.tamu.edu`, an ORSEE database (Greiner, 2015). All sessions were run in the Texas A&M Economic Research Laboratory using z-Tree (Fischbacher, 2007).

At the end of every session, each subject received the time average of their instantaneous payoff plus a five dollar show-up payment. Subject earnings averaged $15.20 in the global-information treatment and $16.09 in the local-information treatment, including the five dollar show-up payment. In equilibrium, average subject earnings would equal $15.00, so subjects received slightly above equilibrium earnings under both treatments. All sessions lasted less than one hour.

the imitative or logit terms (see Table 3).

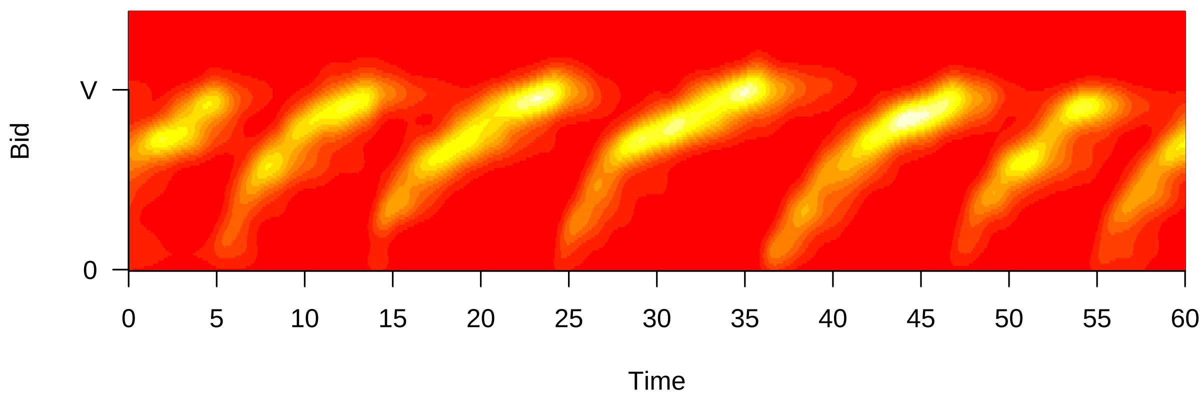

Figure 4: Changes in the distribution of bidding behavior over time in a population of 30 simulated agents under adaptive dynamics obtained from a nonparametric conditional density estimator with a bid bandwidth of 0.5 and a time bandwidth of 0.3 seconds

4 Hypotheses

The game utilized in our experiment is not evolutionary stable (see section 2.2.1), so adaptive models predict persistent disequilibrium rather than convergence to equilibrium. Figure 4 shows a heat map illustrating the changes in the predicted distribution of bidding behavior over time in our experiment for a population of 30 simulated agents. We employ the adaptive dynamics in equation 8 where $\alpha = 3$, $\beta = 0.3$, and $\gamma = 0.3$. In contrast to the static predictions of Nash and Quantal Response Equilibrium model (see section 2.1), the adaptive model predicts persistent bidding cycles.[7] This theoretical prediction results in the following hypothesis.

Hypothesis 1 *Subjects will exhibit persistent bidding cycles under both the local information treatment and the social information treatment.*

Throughout these disequilibrium cycles, adaptive models of imitation and optimization predict very different behavioral dynamics. Figure 5 provides an example of the predicted

[7] The intuition for bidding cycles is as follows. When bids are sufficiently low agents can benefit by slightly outbidding their competitors, so competition gradually drives bids upwards. As bids gradually increase towards the value of the prize, average profits decrease. When profits became sufficiently low, agents can effectively opt-out of the auction by bidding close to zero, thus reinitializing the bidding cycle.

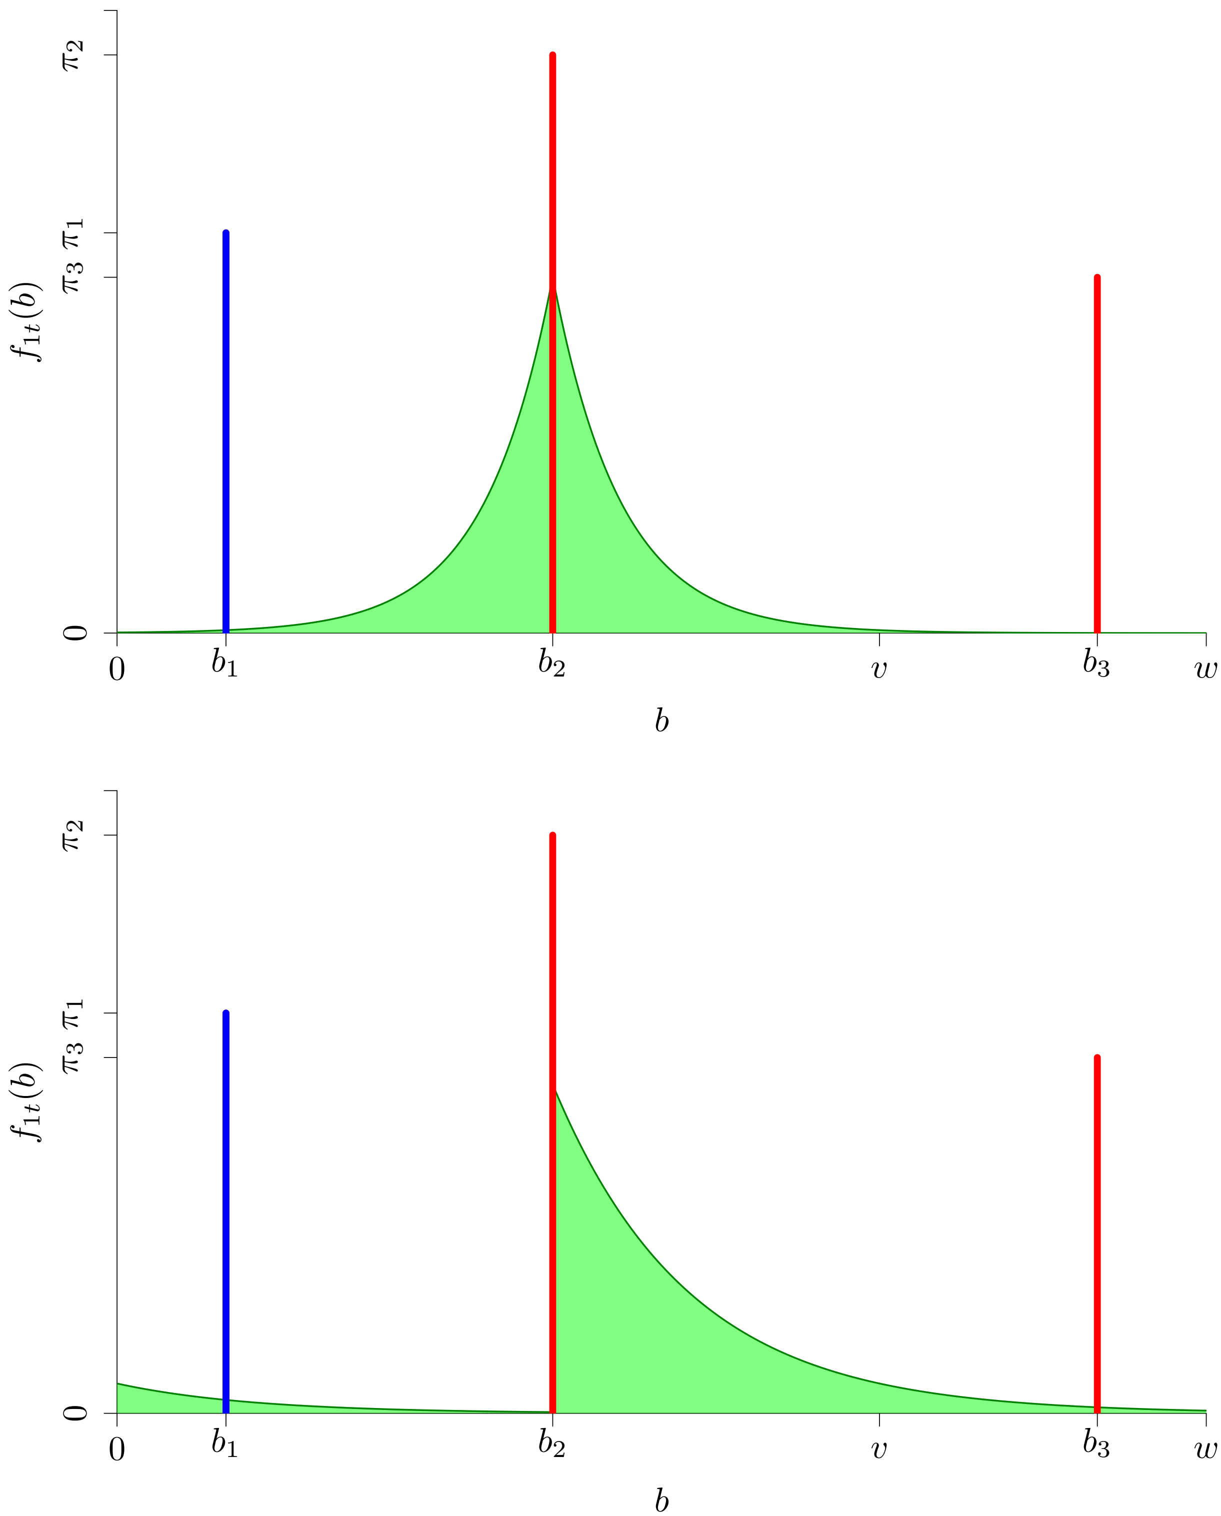

Figure 5: The predicted probability density of a new bid selected by player 1 at time t. The top panel illustrates the predicted density under imitative models and the lower panel illustrates the predicted density under myopic optimization models.

probability density of a new bid selected by player 1 at time t. The upper figure illustrates the predicted density under noisy imitative models and the lower figure illustrates the predicted density under noisy myopic optimization models. The horizontal position of each vertical line indicates the current bid of one player and the height of the line indicates the current payoff to this player. The shaded area under the curve indicates the probability density for a new bid selected by player 1.

Under the bidding profile depicted in figure 5, player 2 is has the highest payoff with bid b_2, so noisy imitation models predict that player 1 is likely to imitate this successful strategy by selecting a new bid that is close to b_2. In contrast, noisy optimization models predict that player 1 is only likely to select bids that are slightly higher than b_2 as selecting a bid slightly lower than b_2 would not improve the payoff to player 1. This sharp contrast between the theoretical predictions from imitation and optimization in this game is part of what allows our experimental design to provide a uniquely powerful test of these adaptive models.

Since our experimental design allows us to disentangle imitation from optimization in observed subject behavior, we next ask how these models will perform under our different information treatments. Adaptive imitation models place a strong informational requirement on agents. As Schlag (2011) notes, "imitation is a simple behavior that has two basic ingredients. One needs to be able to observe what others have done and one needs to be capable of doing what they have done." Hence adaptive imitation models require agents to observe the behavior of their peers. In contrast, adaptive optimization models describe agents who attempt to directly maximize their own payoff, so adaptive optimization models require agents only to observe their own payoffs.

Our local information treatment provides each subject with information regarding their own bids and payoffs, so we hypothesize that noisy optimization behavior (depicted in the lower panel of figure 5) will be observed in that treatment. In contrast, information regarding the bids and payoff of others is only provided in our global information treatment, so we hypothesize that noisy imitative behavior (depicted in the top panel of figure 5) will be observed in that treatment. Moreover, our global information treatment is designed to make implementing adaptive imitation as easy as possible; subjects only need to click on the highest bar on a computer screen (see figure 3 for a depiction of the interface) to implement adaptive imitation. In contrast, implementation of adaptive optimization is more computationally demanding for subjects, since it requires them to compute counterfactual payoffs from their information regarding the bids of others. Accordingly, we hypothesize that

Hypothesis 2 *Imitative models will outperform optimization models in the social information treatment, but not in the local information treatment.*

Finally, a growing literature (e.g., Merlo and Schotter, 2003; Armantier, 2004; Cardella, 2012; Brown et al., 2009) suggests that the provision of social information to economic agents helps agents to behave more rationally and come closer to the predictions of traditional economic theory and those of Nash equilibrium. For this reason, we speculate that the additional information provided by our social information treatment may help subjects to learn their way out of disequilibrium cycles and behave more consistently with the theoretical predictions of Nash equilibrium. This reasoning leads to our final hypothesis.

Hypothesis 3 *Behavior in the social information treatment will exhibit greater stability and greater consistency with Nash equilibrium.*

5 Results

Table 1 provides summary statistics for both the local information and global information treatments and a comparison with the equilibrium predictions. Recall that the Nash equilibrium of this game predicts that subjects will employ a mixed strategy with bids distributed according to the probability density function described in Section 2.1.1. On average, subjects in both of our treatments exhibit lower bidding than the equilibrium prediction. Consequently, the average earnings in both treatments are higher than the equilibrium prediction. In both treatments, we also observe instances of dominated bidding—bids above 7—which are never predicted to occur in equilibrium. Consistent with hypothesis 1, we do not observe a convergence to equilibrium in either treatment; the last 10 minutes of the experiment are not noticeably closer to equilibrium play than the first 10 minutes of the experiment. In general, there are only minor differences between the first and last 10 minutes of the ex-

	Private Information			Global Information			Equilibrium Prediction
	initial 10 minutes	last 10 minutes	overall	initial 10 minutes	last 10 minutes	overall	
mean bid	3.73	3.54	3.57	4.57	4.31	4.47	4.67
bids above 7[a]	1.95%	0.45%	1.22%	5.57%	4.58%	6.96%	0.00%
minimum bids[b]	1.63%	3.77%	3.03%	0.72%	0.81%	0.90%	0.00%[c]
mean earnings	10.94	11.13	11.11	10.20	10.36	10.21	10.00

a. In this game, bids above 7 are always dominated by bidding 0.

b. The minimum bid is 0.

c. A bid of 0 is in the support of the mixed equilibrium strategy. However, the predicted occurrence of such bids by the equilibrium model is 0, because the strategy space is continuous.

Table 1: Summary Statistics for Bids and Earnings in Local Information Treatment, Global Information Treatment, and Equilibrium Predictions. The treatment statistics include groupings by the first 10 minutes and last 10 minutes to provide more detail about initial and final play.

periment. Of those differences that exist (e.g., average bid decreasing, earnings increasing), most are moving away, rather than toward, equilibrium predictions.

Result 1 *Both average payoffs and average bids differ significantly across treatments.*

i. Subjects bids are higher and closer to the Nash equilibrium predictions in the global-information treatment.

ii. Payoffs in the global-information treatment are lower and closer to equilibrium than those in the local-information treatment.

Bids in the global-information treatment are significantly higher than those in the local information treatment. Table 1 shows that the mean bid in the global information treatment is $0.90 higher than the mean bid in the local information treatment, so the former is closer to the equilibrium prediction. Subjects in the global information treatment are also more likely to select dominated bids above 7 and less likely to make 0 bids than those in the local information treatment. Figure 6 provides nonparametric estimates of the aggregate bid density under each treatment, showing that bids in the global information treatment are generally larger than those in the local information treatment. A non-parametric Kolmogorov-Smirnov test finds the empirical bid distributions to be significantly different ($p < 0.01$).

Figure 6 illustrates nonparametric estimates of the aggregate bid density under each treatment alongside the symmetric Nash equilibrium density function. As Table 1 implies, both

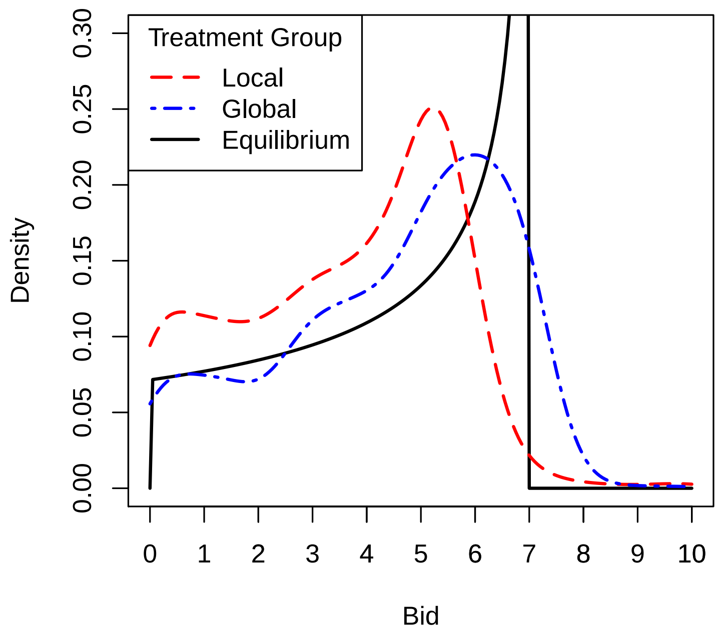

Figure 6: Empirical Bid Distributions. The density function is estimated using the local constant kernel density estimator of Rosenblatt et al. (1956) and Parzen (1962) with a normal kernel and a bandwidth of 0.5.

of the observed bid distributions are generally lower than the equilibrium distribution, but with longer right tails. A non-parametric Kolmogorov-Smirnov test finds both empirically observed bid-distributions to be significantly different from the Nash equilibrium distribution ($p < 0.01$). Moreover, neither bid distribution is consistent with logit quantal response equilibrium. A Kolmogorov-Smirnov test finds both the local-information-treatment bid distribution and the global-information-treatment bid distribution to be significantly different from their corresponding maximum-likelihood logit quantal response equilibrium predictions ($p < 0.01$; local information: $\eta = 0.792$, global information: $\eta = 0.505$).

Payoffs also differ significantly across treatments. Table 1 shows that mean payoffs in the local information treatment are \$0.90 higher than the global information treatment. Subjects in the global information treatment earned an average of \$15.21 while subjects in the local information treatment, subjects earned an average of \$16.11; significantly higher earnings at the one percent level. A non-parametric Kolmogorov-Smirnov test finds the empirical payoff

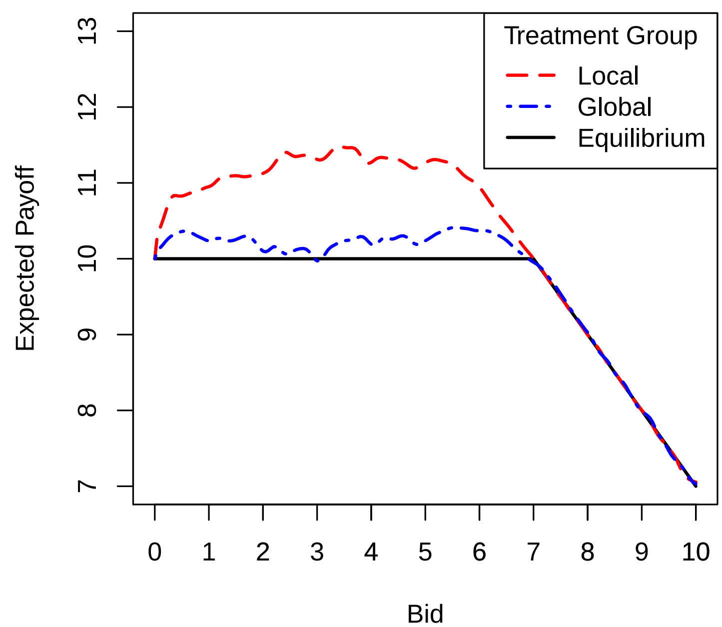

Figure 7: Empirical Expected Payoff Functions under Alternate Treatments. The expected payoff function is estimated using the local constant kernel regerssion estimator of Nadaraya (1964) and Watson (1964) with a normal kernel and a bandwidth of 0.5

distributions to be significantly different ($p < 0.01$).

In equilibrium, every bid between zero and the value of the prize should yield the same expected payoff since rational agents must be indifferent between pure strategies over which they mix. Figure 7 shows that both treatments violate this indifference property. However, this violation is much more severe in the local information treatment than in the global information treatment, suggesting that subjects in the global information treatment are more precisely maximizing their payoffs. A non-parametric Kolmogorov-Smirnov test finds both empirically observed earnings distributions to be significantly different from the Nash equilibrium distribution ($p < 0.01$). Similarly, the empirical earnings distributions are also inconsistent with quantal response equilibrium. A Kolmogorov-Smirnov test finds both the local-information-treatment earnings distribution and the global-information-treatment earnings distribution to be significantly different from their corresponding maximum-likelihood, quantal response equilibrium predictions ($p < 0.01$; local information: $\eta = 0.792$, global information: $\eta = 0.505$).

Result 2 *Throughout the* 40 *minute session, subject behavior in both treatments is characterized by a state of disequilibrium, resembling neither the Nash-equilibrium nor the logit quantal response equilibrium. There is no convergence to equilibrium; rather behavior in both treatments is characterized by persistent cycling.*

Consistent with hypothesis 1, the bidding behavior observed in each treatment is characterized by persistent, identifiable, disequilibrium cycling. Figures 8 and 9 illustrate “heat maps” for the same periods as Figures 10(a-f). Moreover, note that the observed cycles in bidding behavior are consistent with the theoretical predictions from adaptive models illustrated by figure 4. In contrast to Hypothesis 3, these cycles are more rapid in the global-information treatment than in the local information treatment, so bidding behavior is actually less stable under the presence of social information.

One quantitative way to analyze the observed bidding behavior is to examine the time series of the mean bid employed by subjects. Figures 10(a-c) illustrate the dynamics of the average bid in the global information treatment for the first, middle, and last minute, respectively, of the session. There is also a strong cyclical pattern to the mean bid in all three phases. Figures 10(d-f) provide the dynamics of the average bid in the local information treatment for the first, middle, and last minute, respectively, of the session. If subjects employ an equilibrium mixed strategy, then future changes in the mean bid should be uncorrelated with past changes in the mean bid. To test this hypothesis, we conduct the Ljung–Box test on the differenced time series of the mean bid. We find that the Ljung–Box test rejects the null hypothesis of uncorrelated changes in the mean bid at the one percent level under both treatments, suggesting that subjects exhibit significant disequilibrium dynamics in both treatments.

Result 3 *Observed bidding dynamics differ significantly across treatments.*

i. Bidding cycles have higher frequency in the global information treatment than in local information treatment.

ii. Bidding dynamics under the global information treatment exhibit far less behavioral noise than under the local information treatment.

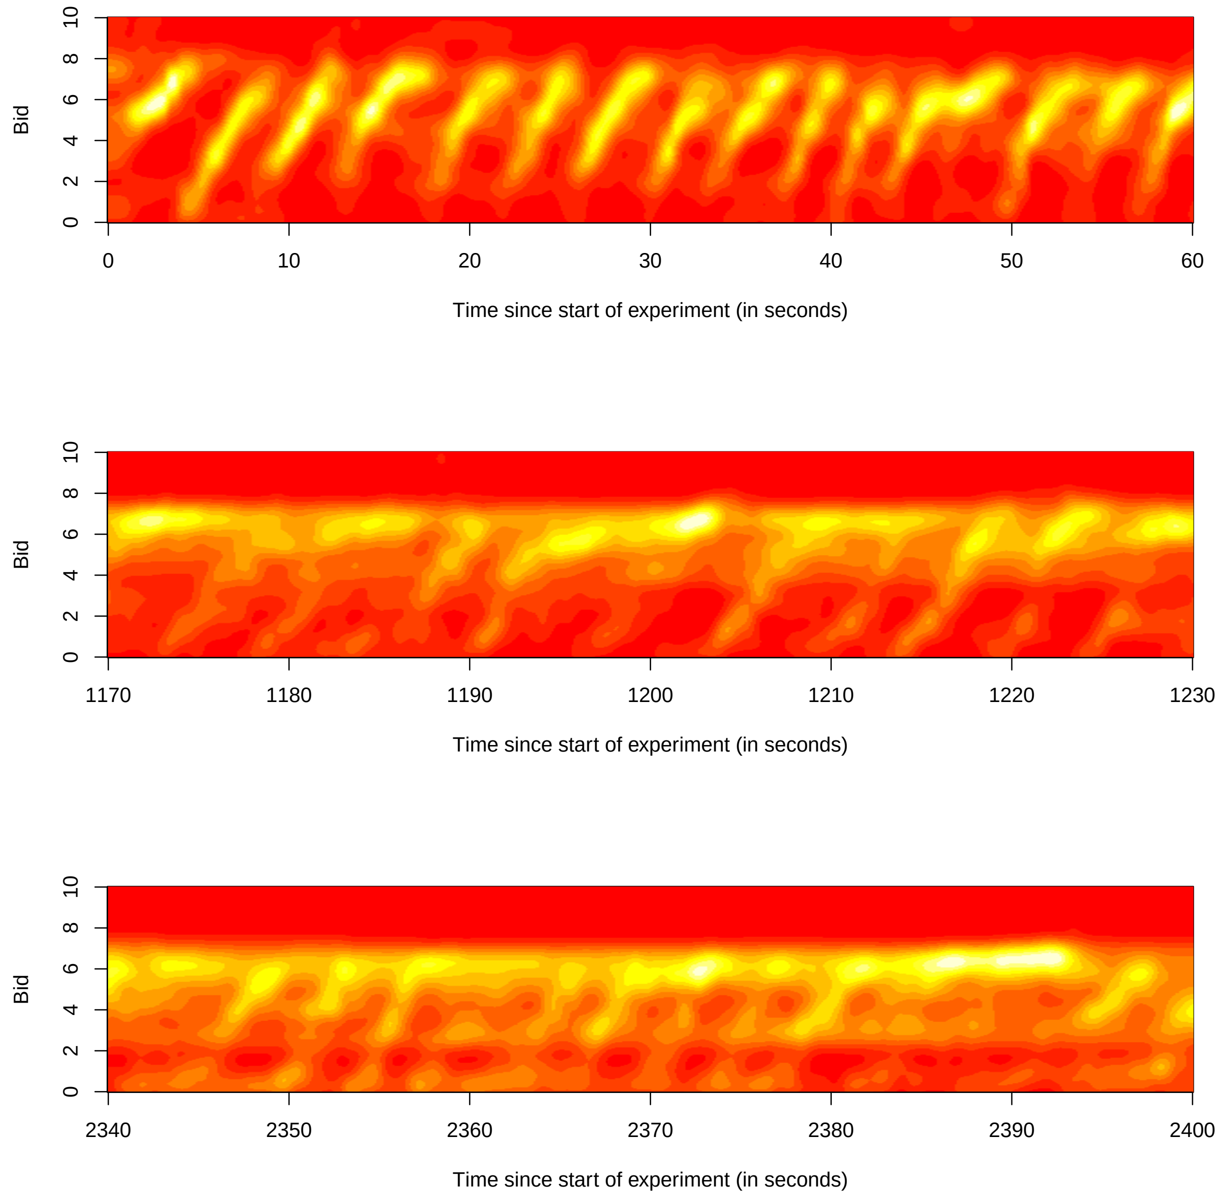

Figure 8: Changes in the empirical distribution of bids over time under global information obtained from a nonparametric conditional density estimator with a bid bandwidth of 0.5 and a time bandwidth of 0.3 seconds. Figures (a-c) depict the global information treatment for the first, median, and last minute, respectively.

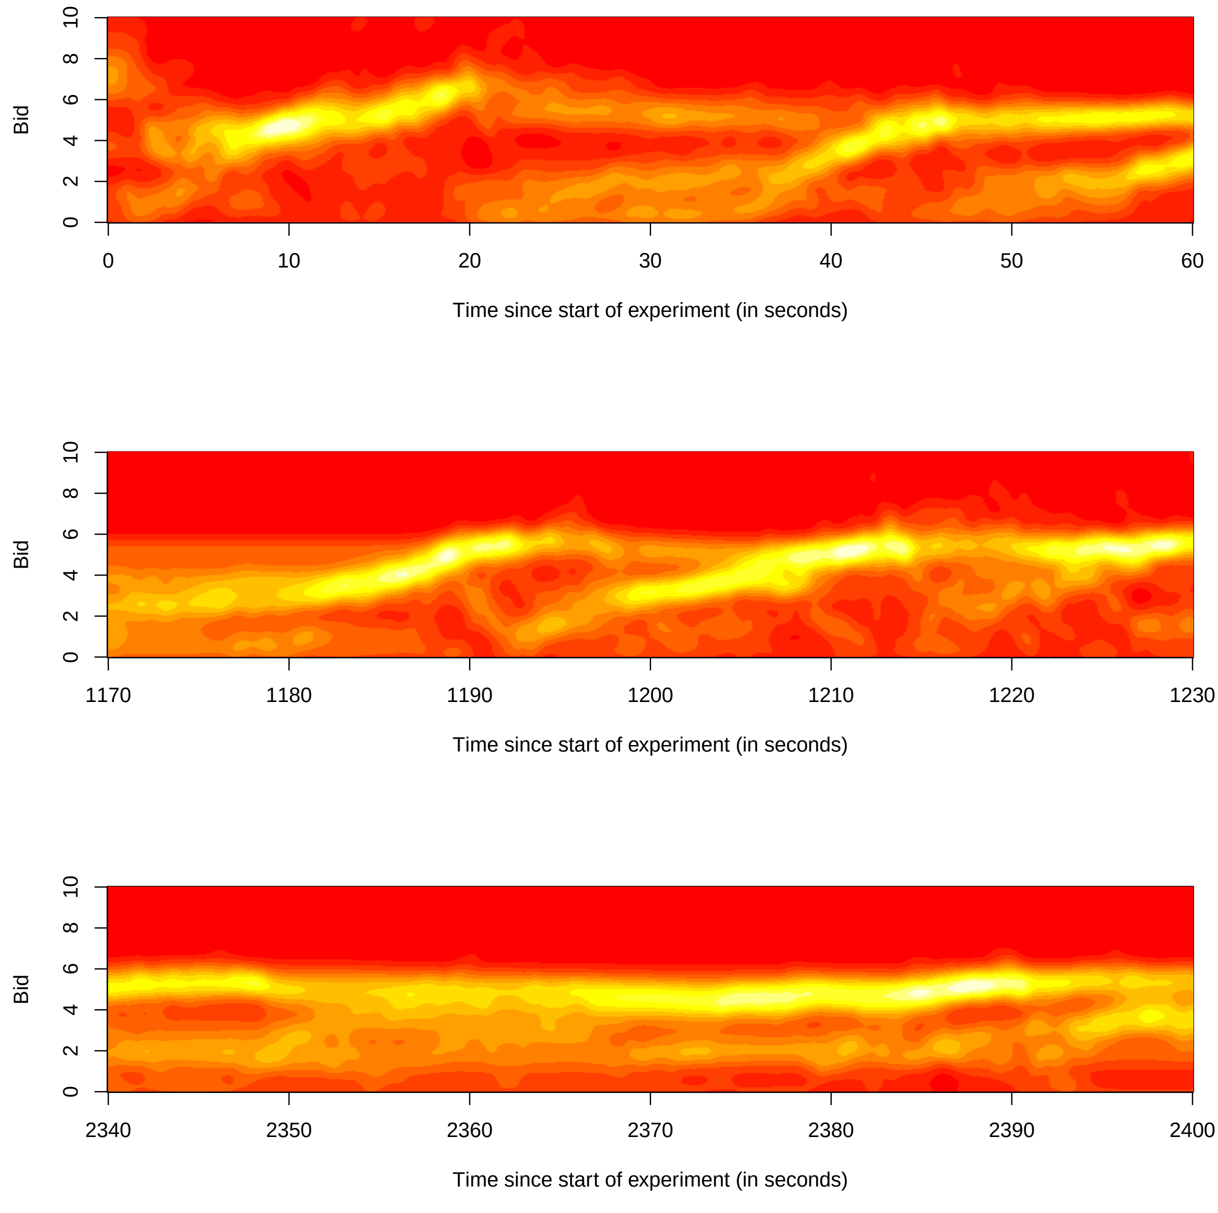

Figure 9: Changes in the empirical distribution of bids over time under local information obtained from a nonparametric conditional density estimator with a bid bandwidth of 0.5 and a time bandwidth of 0.3 seconds. Figures (a-c) depict the local information treatment for the first, median, and last minute, respectively.

Figure 10: Movement of the average bid in the first (top row), median (middle row), and last (bottom row) minutes of the global and local information treatments. Figures (a, c, e) (left) show the global information treatment. Figures (b, d, f) (right) show the local information treatment.

The cycles observed in the aggregate bid data also differ across treatments. Figures 8 and 9 in the global information treatment are characterized by frequent cycles that appear to be about 5 seconds in length. The cycles are more noisy in Figures 10 (b,d,f) in the local information treatment and the distance from the peak of one cycle to the next can be as large as 20 seconds. The heat maps in Figures 8 and 9 also confirm these results. While cyclical patterns repeat about every 5 seconds in the global information treatment, they repeat about every 20-40 seconds in the private information treatment. Consistent with these results, the maximum likelihood estimates reported below indicate greater precision and less autocorrelation in the global information treatment, suggesting that the underlying adaptive processes are significantly different across treatments.

Result 4 *Behavior in both treatments is far more consistent with optimization than imitation.*

It is useful to examine which factors best explain bidding behavior in the observed data. To that end, we estimate adaptive models of both noisy optimization and noisy imitation (see Section 2.2.2 for a full description of these models). Each model provides a continuous probability distribution that gives the likelihood $f_{it}(b)$ that a given bid b will be selected by a subject i at time t, based on a single precision parameter β (see equations 6 and 7 and figure 5 for more details). For a given β, the total likelihood for each of these models is the product of all observed $f_{i,t}(b_{i,t})$. The β^* that maximizes this likelihood function is the maximum likelihood precision parameter. Table 2 provides results for both models under both treatments of the experiment.

There is a clear ranking of these single parameter models in terms of how they explain the data. In contrast to hypothesis 2, optimization dynamics are a better predictor of disequilibrium subject behavior than noisy imitation dynamics under both treatments. Both models have more explanatory power in the global information treatment than in local information treatment, which is consistent with our finding of greater behavioral noise in the local-information treatment than the global-information treatment (see Result 3 for details).

After examining simple one-parameter models separately, it is desirable to combine the noisy optimization, and noisy imitation models in a combined model. We also consider

	Global Information		Local Information	
	Optimization	Imitation	Optimization	Imitation
precision parameter (β)	1.54 (0.15)	0.60 (0.07)	0.68 (0.05)	0.33 (0.12)
observations	32634	32634	27225	27225
total log-likelihood	-58808.48	-66379.35	-57155.04	-58812.47

Table 2: Maximum-Likelihood Models of Noisy optimization and Imitation Dynamics, Global-Information and Local-Information Treatments. The noisy optimzation model outperforms the noisy imitation-response model. Both models perform better in the global information treatment than in the local information treatment. All parameters are significant at the 1% level. Standard errors are obtained via subject clustered bootstrap estimation.

models that include an autocorrelation term accounting for the tendency of subjects to select new bids close to their previous bid. This term is especially relevant in the local information treatment where subjects are unable to observe the bids of others, and hence, tend to employ a trial and error strategy.[8]

Table 3 provides parameter estimates for the combined model in the local-information and global-information treatments. In the global-information treatment, subject bidding behavior is primarily driven by payoff incentives following the noisy optimization dynamic. In addition, there is also some degree to which individuals tend to select bids close to their own previously used bid. Under local information, bidding behavior is largely driven by autocorrelation with previously selected bids, with the payoffs under the noisy imitation dynamic as a secondary factor.

The difference in the explanatory power of the noisy optimization dynamic across treatments is not surprising. Subjects have the ability to directly maximize their payoff only when they have information regarding the bids employed by others. Thus, in the global-information treatment, they can directly respond to their payoff incentives. In the local-information treatment, they only receive information about the payoff they earn from the strategy they currently employ. Without further information about the strategies employed by others, subjects cannot easily determine how their payoff would change if they were to

[8]To make this point more salient, if the bid autocorrelation term were part of our one-parameter model comparison in Table 2, it would provide the greatest explanatory power in the local-information treatment. In the global information treatment, it would still outperform the imitative response model, but not the optimization response model.

	Global Information		Local Information	
logit (α)	1.37	1.48	0.56	0.48
(payoffs)	(0.11)	(0.14)	(0.07)	(0.05)
previous bids (β)	-	0.56	-	1.39
(subject specific)		(0.07)		(0.11)
imitation response (γ)	0.12	0.07	0.10	0.08
(the highest earning bid)	(0.03)	(0.03)	(0.02)	(0.02)
observations	32634	32634	27225	27225
log-likelihood	-58345.36	-49805.06	-56954.91	-33030.53
mean log-likelihood	-1.79	-1.52	-2.09	-1.21
typical bid likelihood	0.17	0.22	0.12	0.30

Table 3: Multiple-Parameter Models of Bidding Dynamics. A multi-parameter model including terms for logit and imitative dynamics is estimated on both the local and global information treatments. An additional specification includes a term for the tendency of subjects to make bids close to their previous bids. All parameters are significant at the 1% level. Standard errors are obtained via subject clustered bootstrap estimation.

adjust their strategy. In this case, it makes sense that subjects would experiment by making trial adjustments and then return to the strategies that provided the highest payoffs. This trial-and-error approach is consistent with the high autocorrelation of current bids and previous bids in the local-information treatment and it also explains how subjects are able to approximately best respond to their opponents' strategies without directly observing them.

Imitation explains very little of the observed behavior in either treatment. In the local information treatment, subjects see neither the payoff nor the strategy of any subject other than themselves, so the lack of imitation is unsurprising, since subjects can not directly implement the imitative model. However, in the global-information treatment, subjects need only to click on the highest bar to perfectly follow the imitative model, but the data indicate that subjects use something more complex than a simple imitation heuristic; they perform noisy myopic payoff optimization.

The autocorrelation in subject bids may explain why the imitative dynamic appeared to have some explanatory power in a one-parameter model. Since bids tend to bunch together, as illustrated in Figure 8, the autocorrelation with a subject's own previous bid produces similar predictions to imitation of the bids employed by others. Consequently, a simple one-parameter imitation model with no autocorrelation parameter can misidentify autocorrelation in a subject's own bids for imitation of others. Results from the combined model presented in Table 3 suggest that much of the explanatory power attributed to imitation under the global informaton treatment actually results from autocorrelation with a subjects own previous bid.

6 Conclusion

This study experimentally investigates dynamic bidding behavior in continuous-time, all-pay auctions. In contrast to previous experimental studies of the all-pay auction, our subjects earned continuous flow payoffs and could adjust their bids asynchronously throughout the experiment. By permitting this type of asynchronous adjustment, we obtain a remarkably fine-grained picture of the empirical bidding behavior, allowing a close examination of behavioral bidding dynamics.

Consistent with theoretical predictions from adaptive models, but in contrast with both Nash and quantal response equilibrium predictions, subjects in our experiment exhibited persistent cyclical bidding behavior. This sustained disequilibrium behavior, along with the markedly discontinuous nature of payoff functions in the all-pay auction, allows us to closely investigate the predictive power of imitative and optimization dynamics. Surprisingly, behavior in the global-information treatment, which provides each subject with the information to easily employ imitative dynamics, is characterized by increased precision of optimization behavior but very little imitative behavior, resulting in higher bids, lower payoffs, and more rapid behavioral cycles.

Our results suggest a general failure of imitative models to adequately describe human cognition in strategic settings. Subjects in the global information treatment could easily imitate the highest performing subject by selecting the highest line on a computer screen. However, instead of merely imitating successful strategies, subjects followed more sophisticated optimization methods, responding to the structure of their payoff incentives. In the local-information treatment, subjects do not have the necessary information to imitate other subjects. In the absence of social information, subjects employ trial-and-error strategies, selecting strategies near those that gave them higher payoffs. Subjects in the global information treatment compete more vigorously, their bidding cycles are far more rapid, and they exhibit far less behavioral noise. As a result, both average bids and average earnings are significantly closer to equilibrium predictions in the global information treatment than in the local information treatment.

While this experiment is primarily concerned with testing theoretical predictions, it also provides some interesting policy implications. In particular, these results suggest that policy makers may want to promote the distribution of social information in strategic environments where effort expenditure has positive externalities, such as patent races or competition for research grants. In contrast, policy makers may want to discourage the distribution of social information in strategic environments where effort expenditure is wasteful or has negative externalities, such as political lobbying or international warfare. Naturally, further research will be needed to verify the extent to which these experimental results carry over to other strategic environments.

References

Alos-Ferrer, C. and Ania, A. B. (2005). The evolutionary stability of perfectly competitive behavior. *Economic Theory*, 26(3):497–516.

Anderson, S. P., Goeree, J. K., and Holt, C. A. (1998). Rent seeking with bounded rationality: An analysis of the all-pay auction. *Journal of Political Economy*, 106(4):828–853.

Armantier, O. (2004). Does observation influence learning? *Games and Economic Behavior*, 46(2):221–239.

Baye, M. R., Kovenock, D., and De Vries, C. G. (1993). Rigging the lobbying process: an application of the all-pay auction. *American Economic Review*, 83(1):289–294.

Brown, A. L., Chua, Z. E., and Camerer, C. F. (2009). Learning and visceral temptation in dynamic saving experiments. *The Quarterly Journal of Economics*, pages 197–231.

Camerer, C. and Hua Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874.

Cardella, E. (2012). Learning to make better strategic decisions. *Journal of Economic Behavior & Organization*, 84(1):382–392.

Cason, T. N., Friedman, D., and Hopkins, E. (2014). Cycles and instability in a rock-paper-scissors population game: A continuous time experiment. *Review of Economic Studies*, 81(1):112–136.

Chatterjee, K., Reiter, J. G., and Nowak, M. A. (2012). Evolutionary dynamics of biological auctions. *Theoretical Population Biology*, 81(1):69–80.

Cheung, Y.-W. and Friedman, D. (1998). A comparison of learning and replicator dynamics using experimental data. *Journal of Economic Behavior & Organization*, 35(3):263–280.

Clark, D. J. and Riis, C. (1998). Competition over more than one prize. *American Economic Review*, 88(1):276–289.

Cournot, A., Bacon, N., and Fisher, I. (1897). *Researches Into the Mathematical Principles of the Theory of Wealth*. Economic classics. Macmillan.

Dechenaux, E., Kovenock, D., and Sheremeta, R. M. (2014). A survey of experimental research on contests, all-pay auctions and tournaments. *Experimental Economics*, 18(4):609–669.

Demichelis, S. and Weibull, J. W. (2008). Language, meaning, and games: A model of communication, coordination, and evolution. *American Economic Review*, 98(4):1292–1311.

Duersch, P., Oechssler, J., and Schipper, B. C. (2012). Unbeatable imitation. *Games and Economic Behavior*, 76(1):88–96.

Ernst, C. and Thöni, C. (2013). Bimodal bidding in experimental all-pay auctions. *Games*, 4(4):608–623.

Feri, F., Meléndez-Jiménez, M. A., Ponti, G., and Vega-Redondo, F. (2011). Error cascades in observational learning: An experiment on the chinos game. *Games and Economic Behavior*, 73(1):136–146.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Friedman, D., Huck, S., Oprea, R., and Weidenholzer, S. (2015). From imitation to collusion: Long-run learning in a low-information environment. *Journal of Economic Theory*, 155:185–205.

Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*, volume 2. MIT press.

Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica*, 59(3):859–867.

Gneezy, U. and Smorodinsky, R. (2006). All-pay auctions: an experimental study. *Journal of Economic Behavior & Organization*, 61(2):255–275.

Golman, R. (2011). Why learning doesnt add up: equilibrium selection with a composition of learning rules. *International Journal of Game Theory*, 40(4):719–733.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125.

Ho, T. H., Camerer, C. F., and Chong, J.-K. (2007). Self-tuning experience weighted attraction learning in games. *Journal of Economic Theory*, 133(1):177–198.

Hodler, R. and Yektaş, H. (2012). All-pay war. *Games and Economic Behavior*, 74(2):526–540.

Hopkins, E. (1999). A note on best response dynamics. *Games and Economic Behavior*, 29(1):138–150.

Lopez, G. (1995). Quantal response equilibria for models of price competition. *Unpublished Ph. D. dissertation, University of Virginia.*

Lugovskyy, V., Puzzello, D., and Tucker, S. (2010). An experimental investigation of overdissipation in the all pay auction. *European Economic Review*, 54(8):974–997.

Marinucci, M. and Vergote, W. (2011). Endogenous network formation in patent contests and its role as a barrier to entry. *Journal of Industrial Economics*, 59(4):529–551.

McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38.

Merlo, A. and Schotter, A. (2003). Learning by not doing: an experimental investigation of observational learning. *Games and Economic Behavior*, 42(1):116–136.

Nadaraya, E. A. (1964). On estimating regression. *Theory of Probability & Its Applications*, 9(1):141–142.

Nash, J. (1951). Non-cooperative games. *Annals of Mathematics*, 54(2):286–295.

Oechssler, J. and Riedel, F. (2001). Evolutionary dynamics on infinite strategy spaces. *Economic Theory*, 17(1):141–162.

Offerman, T., Potters, J., and Sonnemans, J. (2002). Imitation and belief learning in an oligopoly experiment. *Review of Economic Studies*, 69(4):973–997.

Oprea, R., Henwood, K., and Friedman, D. (2011). Separating the hawks from the doves: Evidence from continuous time laboratory games. *Journal of Economic Theory*, 146(6):2206–2225.

Parayre, R. and Hurry, D. (2001). Corporate investment and strategic stability in hypercompetition. *Managerial and Decision Economics*, 22(4-5):281–298.

Parzen, E. (1962). On estimation of a probability density function and mode. *Annals of Mathematical Statistics*, 33(3):1065–1076.

Rosenblatt, M. et al. (1956). Remarks on some nonparametric estimates of a density function. *Annals of Mathematical Statistics*, 27(3):832–837.

Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT press.

Schlag, K. (2011). Imitation and social learning. In Seel, N., editor, *Encyclopedia of the Sciences of Learning*, pages 1489–1493. Springer Science & Business Media.

Smith, J. M. and Price, G. (1973). The logic of animal conflict. *Nature*, 246:15–18.

Smith, M. J. (1984). The stability of a dynamic model of traffic assignment-an application of a method of lyapunov. *Transportation Science*, 18(3):245–252.

Taylor, P. D. and Jonker, L. B. (1978). Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1):145–156.

Vega-Redondo, F. (1997). The evolution of walrasian behavior. *Econometrica*, 65(2):375–384.

Watson, G. S. (1964). Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 359–372.

Weibull, J. (1997). *Evolutionary Game Theory*. MIT Press.

A Mathematical Appendix

A.1 Nash Equilibrium Derivation

Consider the following auction with three bidders and two prizes. Each player starts with an endowment w selects her bid from the closed interval $[0, w]$. After all three bids have been selected, the top two bidders each receive a prize with value $v < w$. However, every player must pay her bid, regardless of whether or not she won a prize. In the case of a tie, the remaining prizes are randomly assigned among the tying players. Accordingly, the payoff function for player i is given by:

$$\pi_i\left(s_i, s_j, s_k\right) = \begin{cases} w - s_i + v & \text{if } s_i > \min\left\{s_j, s_k\right\} \\ w - s_i + 2v/3 & \text{if } s_i = s_j = s_k \\ w - s_i + v/2 & \text{if } s_i = \min\left\{s_j, s_k\right\} < \max\left\{s_j, s_k\right\} \\ w - s_i & \text{otherwise} \end{cases}$$

Suppose that there exists a continuous symmetric mixed strategy Nash equilibrium with support over the closed interval $[0, v]$. Let $F(z) = P(b_j < z)$ denote the corresponding cumulative distribution function. Let W_i denote the event that bidder i wins and receives and item. Let L_i denote the event that bidder i loses and does not receive and item. If bidder j and bidder k follow this Nash equilibrium mixed strategy, then the probability that bidder i loses the auction is given by:

$$\begin{aligned} P(L_i) &= P(b_i < b_j \text{ and } b_i < b_k) \\ &= P(b_i < b_j) P(b_i < b_k) = P(b_i < b_j)^2 \\ &= [1 - P(b_j < b_i)]^2 = [1 - F(b_i)]^2 \\ &= 1 - 2F(b_i) + F(b_i)^2 \end{aligned}$$

Now the probability that bidder i wins the auction is given by:

$$\begin{aligned} P(W_i) &= 1 - P(L_i) \\ &= 1 - \left[1 - 2F(b_i) + F(b_i)^2\right] \\ &= 2F(b_i) - F(b_i)^2 \end{aligned}$$

Hence bidder i's expected payoff, conditional on her bid, is given by:

$$\begin{aligned} \pi_i(b_i) &= w + vP(W_i) - b_i \\ &= w + v\left[2F(b_i) - F(b_i)^2\right] - b_i \\ &= w + 2vF(b_i) - vF(b_i)^2 - b_i \end{aligned}$$

If the mixed strategy F is a best response for agent i, then she must be indifferent between all of the bids in the support of F. Hence all of the bids in the closed interval $[0, v]$ must yield the same expected payoff for bidder i. Moreover, since bidding zero will certainly yield an expected payoff of zero, every bid between zero and v must yield an expected payoff of zero. Accordingly, we can write:

$$\begin{aligned}
\pi_i(b_i) &= 0 \qquad \text{for all } b_i \in [0,v] \\
2vF(b_i) - vF(b_i)^2 - b_i &= 0 \\
-b_i &= -2vF(b_i) + vF(b_i)^2 \\
-\frac{b_i}{v} &= -2F(b_i) + F(b_i)^2 \\
1 - \frac{b_i}{v} &= 1 - 2F(b_i) + F(b_i)^2 \\
\sqrt{1 - \frac{b_i}{v}} &= 1 - F(b_i) \\
F(b_i) &= 1 - \sqrt{1 - \frac{b_i}{v}} \qquad \text{for all } b_i \in [0,v]
\end{aligned}$$

Differentiating this cumulative distribution function obtains the Nash equilibrium probability density function:

$$f(b_i) = \frac{1}{2v}\left(1 - \frac{b_i}{v}\right)^{-1/2} \qquad \text{for all } b_i \in [0,v]$$

A.2 Logit Quantal Response Equilibrium Derivation

Consider the following auction with three bidders and two prizes. Each player starts with an endowment w selects her bid from the closed interval $[0, w]$. After all three bids have been selected, the top two bidders each receive a prize with value $v < w$. However, every player must pay her bid, regardless of whether or not she won a prize. In the case of a tie, the remaining prizes are randomly assigned among the tying players. Accordingly, the payoff function for player i is given by:

$$\pi_i(s_i, s_j, s_k) = \begin{cases} w - s_i + v & \text{if } s_i > \min\{s_j, s_k\} \\ w - s_i + 2v/3 & \text{if } s_i = s_j = s_k \\ w - s_i + v/2 & \text{if } s_i = \min\{s_j, s_k\} < \max\{s_j, s_k\} \\ w - s_i & \text{otherwise} \end{cases}$$

Suppose that there exists a continuous symmetric logit quantal response equilibrium with support over the closed interval $[0, w]$. Let $F(z) = P(b_j < z)$ denote the corresponding cumulative distribution function. Let W_i denote the event that bidder i wins and receives and item. Let L_i denote the event that bidder i loses and does not receive and item. If bidder j and bidder k follow this mixed strategy, then the probability that bidder i loses the auction is given by:

$$
\begin{aligned}
P(L_i) &= P(b_i < b_j \text{ and } b_i < b_k) \\
&= P(b_i < b_j) P(b_i < b_k) = P(b_i < b_j)^2 \\
&= [1 - P(b_j < b_i)]^2 = [1 - F(b_i)]^2 \\
&= 1 - 2F(b_i) + F(b_i)^2
\end{aligned}
$$

Accordingly, the probability that bidder i wins the auction is given by:

$$
\begin{aligned}
P(W_i) &= 1 - P(L_i) \\
&= 1 - \left[1 - 2F(b_i) + F(b_i)^2\right] \\
&= 2F(b_i) - F(b_i)^2
\end{aligned}
$$

Hence bidder i's expected payoff, conditional on her bid, is given by:

$$
\begin{aligned}
\pi_i(b_i) &= w + vP(W_i) - b_i \\
&= w + v\left[2F(b_i) - F(b_i)^2\right] - b_i \\
&= w + 2vF(b_i) - vF(b_i)^2 - b_i
\end{aligned}
$$

Under a logit quantal response equilibrium, agents do not always select their best response, but they are more likely to select bids that yield higher payoffs. Here the level of behavioral noise is indexed by the parameter η. As η approaches infinity, the logit quantal response equilibrium approaches uniformly random behavior. As η approaches zero, the logit quantal response equilibrium approximates a Nash equilibrium. Formally, we can write:

$$\begin{aligned} f(b) &= \frac{\exp(\eta\pi_i(b))}{\int_0^w \exp(\eta\pi_i(x))\,dx} \\ f(b) &= \frac{\exp\left(\eta\left(2vF(b) - vF(b)^2 - b\right)\right)}{C_0} \\ C_0 f(b) &= \exp\left(2\eta vF(b) - \eta vF(b)^2 - \eta b\right) \\ C_0\frac{dF}{db} &= \exp\left(2\eta vF(b) - \eta vF(b)^2 - \eta b\right) \end{aligned}$$

Integrating both sides of this differential equation obtains

$$\begin{aligned} C_0\int \exp\left(\eta vF^2 - 2\eta vF\right)dF &= \exp(-\eta b)\,db \\ \eta C_0\int \exp\left(\eta vF^2 - 2\eta vF\right)dF &= \frac{1}{\eta} - \frac{1}{\eta}\exp(-\eta b) \\ C_1\exp(\eta v)\int \exp\left(\eta v(F-1)^2\right)dF &= \frac{1}{\eta} - \frac{1}{\eta}\exp(-\eta b) \\ \int \exp\left(\eta v(F-1)^2\right)dF &= C_3 - C_4\exp(-\eta b) \end{aligned}$$

We can solve for the cumulative distribution function F in terms of the the imaginary error function by introducing the function $G(x) = \int_0^x \exp(u^2)\,du = \frac{\sqrt{\pi}}{2}\text{erfi}(x)$.

$$\begin{aligned} G\left(\sqrt{\eta v}(F-1)\right) &= C_3 - C_4\exp(-\eta b) \\ \sqrt{\eta v}(F-1) &= G^{-1}\left(C_3 - C_4\exp(-\eta b)\right) \\ F-1 &= \frac{1}{\sqrt{\eta v}}G^{-1}\left(C_3 - C_4\exp(-\eta b)\right) \\ F(b) &= 1 - \frac{1}{\sqrt{\eta v}}G^{-1}\left(C_3 - C_4\exp(-\eta b)\right) \end{aligned}$$

Now since bids are restricted to be non-negative, we have

$$\begin{aligned} F(0) &= 0 \\ 1-\frac{1}{\sqrt{\eta v}}G^{-1}\left(C_3 - C_4\exp(0)\right) &= 0 \\ \frac{1}{\sqrt{\eta v}}G^{-1}\left(C_3 - C_4\right) &= 1 \\ G^{-1}\left(C_3 - C_4\right) &= \sqrt{\eta v} \\ C_3 - C_4 &= G\left(\sqrt{\eta v}\right) \\ C_3 &= C_4 + G\left(\sqrt{\eta v}\right) \end{aligned}$$

Similarly, since bids cannot exceed the endowment w, we have

$$\begin{aligned} F(w) &= 1 \\ 1-\frac{1}{\sqrt{\eta v}}G^{-1}\left(C_3 - C_4\exp(-\eta w)\right) &= 1 \\ \frac{1}{\sqrt{\eta v}}G^{-1}\left(C_3 - C_4\exp(-\eta w)\right) &= 0 \\ G^{-1}\left(C_3 - C_4\exp(-\eta w)\right) &= 0 \\ C_3 - C_4\exp(-\eta w) &= G(0) = 0 \\ C_4 + G\left(\sqrt{\eta v}\right) - C_4\exp(-\eta w) &= 0 \qquad \text{since } C_3 = C_4 + G\left(\sqrt{\eta v}\right) \\ G\left(\sqrt{\eta v}\right) &= C_4(\exp(-\eta w) - 1) \\ C_4 &= \frac{G\left(\sqrt{\eta v}\right)}{\exp(-\eta w) - 1} \end{aligned}$$

We can use these solutions for C_3 and C_4 to obtain a closed form solution for the cumulative distribution function F

$$\begin{aligned} F(b) &= 1-\frac{1}{\sqrt{\eta v}}G^{-1}\left(C_3 - C_4\exp(-\eta b)\right) \\ F(b) &= 1-\frac{1}{\sqrt{\eta v}}G^{-1}\left(C_4 + G\left(\sqrt{\eta v}\right) - C_4\exp(-\eta b)\right) \qquad \text{since } C_3 = C_4 + G\left(\sqrt{\eta v}\right) \\ F(b) &= 1-\frac{1}{\sqrt{\eta v}}G^{-1}\left(G\left(\sqrt{\eta v}\right)\left[1-\frac{1-\exp(-\eta b)}{1-\exp(-\eta w)}\right]\right) \qquad \text{since } C_4 = \frac{G\left(\sqrt{\eta v}\right)}{\exp(-\eta w) - 1} \end{aligned}$$

Differentiating the cumulative distribution function F obtains the corresponding probability density function

$$\begin{aligned}
f(b) &= F'(b) = -\frac{1}{\sqrt{\eta v}}\frac{\partial}{\partial b}\left[G^{-1}\left(G\left(\sqrt{\eta v}\right)\left[1-\frac{1-\exp(-\eta b)}{1-\exp(-\eta w)}\right]\right)\right] \\
f(b) &= -\frac{1}{\sqrt{\eta v}}\frac{\partial}{\partial b}\left[G^{-1}\left(H(b)\right)\right] \\
f(b) &= -\frac{1}{\sqrt{\eta v}}\frac{\partial G^{-1}(H(b))}{\partial H(b)}H'(b) \\
f(b) &= -\frac{1}{\sqrt{\eta v}}G'(G^{-1}(H(b)))^{-1}H'(b) \\
f(b) &= -\frac{1}{\sqrt{\eta v}}\exp(G^{-1}(H(b))^2)^{-1}H'(b) \qquad \text{since } G'(x)=\exp\left(x^2\right) \\
f(b) &= -\frac{1}{\sqrt{\eta v}}\exp(-G^{-1}(H(b))^2)H'(b) \\
f(b) &= -\frac{1}{\sqrt{\eta v}}\exp(-G^{-1}(H(b))^2)H'(b) \\
f(b) &= \frac{\eta G\left(\sqrt{\eta v}\right)\exp(-\eta b)}{\sqrt{\eta v}\left[1-\exp(-\eta w)\right]}\exp(-G^{-1}(H(b))^2) \qquad \text{since } H'(b) = -\frac{\eta G\left(\sqrt{\eta v}\right)\exp(-\eta b)}{\left[1-\exp(-\eta w)\right]} \\
f(b) &= \frac{\eta G\left(\sqrt{\eta v}\right)\exp(-\eta b)}{\sqrt{\eta v}\left[1-\exp(-\eta w)\right]}\exp\left(-G^{-1}\left(G\left(\sqrt{\eta v}\right)\left[1-\frac{1-\exp(-\eta b)}{1-\exp(-\eta w)}\right]\right)^2\right)
\end{aligned}$$

A.3 Evolutionary Instability of the Nash Equilibrium

Intuitively, a strategy is evolutionarily stable if it induces a self-enforcing convention. In other words, a strategy x is evolutionarily stable if no other strategy y can invade it when the entire population initially employs strategy x. More formally, in a in a symmetric normal form game, a strategy x is evolutionarily stable if there exists some $C \in (0,1)$ such that for all $\varepsilon \in (0, C)$ and for any other strategy y

$$\pi\left(\, x \mid \epsilon y + (1-\epsilon)\, x \,\right) > \pi\left(\, y \mid \epsilon y + (1-\epsilon)\, x \,\right) \tag{10}$$

Thus, if x is evolutionarily stable and a sufficiently small proportion of the population deviates to an alternate strategy y, then agents who employ x will earn a strictly higher payoff than agents who employ y.

The Nash equilibrium strategy for the all-pay auction is not evolutionarily stable. To see why, suppose that a small proportion ϵ of the population deviates from the Nash equilibrium strategy x to an alternate strategy y under which agents always bid the full value of the prize. Since the support of the equilibrium bid distribution is given by the closed interval $[0, v]$, agents who employ the invading strategy y will win the prize with probability one whenever they are matched against an agent who employs the equilibrium bidding strategy. So the expected payoff to an agent who deviates to strategy y is given by

$$
\begin{aligned}
\pi\,(\ y \mid \varepsilon y + (1-\varepsilon)x\ ) &= \varepsilon^2 \pi_1\,(y,y,y) + 2\varepsilon\,(1-\varepsilon)\,\pi_1\,(y,y,x) + (1-\varepsilon)^2\,\pi_1\,(y,x,x) \\
\pi\,(\ y \mid \varepsilon y + (1-\varepsilon)x\ ) &= \varepsilon^2 \pi_1\,(y,y,y) \qquad \text{since } \pi\,(y,y,x) = \pi_1\,(y,x,x) = 0 \\
\pi\,(\ y \mid \varepsilon y + (1-\varepsilon)x\ ) &= -\frac{\varepsilon^2 v}{3} \qquad \text{since } \pi_1\,(y,y,y) = -\frac{v}{3}
\end{aligned}
$$

On the other hand, the expected payoff to an agent who employs the equilibrium mixed strategy is given by

$$
\begin{aligned}
\pi\,(\ x \mid \varepsilon y + (1-\varepsilon)x\ ) &= \varepsilon^2 \pi_1\,(x,y,y) + 2\varepsilon\,(1-\varepsilon)\,\pi_1\,(x,y,x) + (1-\varepsilon)^2\,\pi_1\,(x,x,x) \\
\pi\,(\ x \mid \varepsilon y + (1-\varepsilon)x\ ) &= \varepsilon^2 \pi_1\,(x,y,y) + 2\varepsilon\,(1-\varepsilon)\,\pi_1\,(x,y,x) \qquad \text{since } \pi_1\,(x,x,x) = 0 \\
\pi\,(\ x \mid \varepsilon y + (1-\varepsilon)x\ ) &< -\varepsilon^2 \pi_1\,(x,y,y) \qquad \text{since } \pi_1\,(x,y,x) < 0 \\
\pi\,(\ x \mid \varepsilon y + (1-\varepsilon)x\ ) &< -\varepsilon^2 E\{bid|x\} \qquad \text{since } \pi_1\,(x,y,y) = -E\{bid|x\} \\
\pi\,(\ x \mid \varepsilon y + (1-\varepsilon)x\ ) &< -\frac{2\varepsilon^2 v}{3} \qquad \text{since } E\{bid|x\} = \frac{2v}{3} \\
\pi\,(\ x \mid \varepsilon y + (1-\varepsilon)x\ ) &< \pi_1\left(y, (\varepsilon y + (1-\varepsilon)x)^2\right) \qquad \text{since } -\frac{2\varepsilon^2 v}{3} < -\frac{\varepsilon^2 v}{3}
\end{aligned}
$$

Thus the invading strategy y earns a higher expected payoff than the equilibrium mixed strategy x, so the equilibrium mixed strategy for the all-pay auction with three bidders and two prizes is not evolutionarily stable. Hence the mixed strategy Nash equilibrium does not induce a self enforcing convention in this all-pay auction. Accordingly, we expect to observe dynamic instability in experimental bidding behavior.