Predicting Behavior in Disequilibrium in Continuous Space and Time: What do Equilibrium, Stability Criteria, and Adaptive Models Get Right?*

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

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

April 6, 2017

Abstract

In some strategic environments, the long-run behavior of boundedly rational agents converges to Nash equilibrium, but in others, behavior may be persistently non-convergent. In this paper, we derive and experimentally test theoretical predictions regarding long-run population-level behavior in all-pay auctions where all but the lowest bidder win a prize. We compare two-prize to one-prize auctions. Consistent with evolutionary stability criteria, we find bidding behavior is more stable in one-prize auctions. Consistent with adaptive models, we observe persistent disequilibrium cycles. These results suggest that stability criteria can help predict whether Nash equilibrium provides a reliable characterization of long-run behavior and that adaptive models can help characterize long-run behavior in unstable strategic environments. As an orthogonal treatment, we also vary the presentation of information given to subjects by focusing on individual payoffs or others' performance. We note subtle changes in disequilibrium dynamics across informational treatments.

*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 three anonymous referees and the attendees of both the 2013 North American and World Meetings of the Economic Science Association.

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

[‡]Email: alexbrown@tamu.edu.

1 Introduction

Equilibrium is a powerful tool for understanding strategic behavior. Even in very complex games, equilibrium predictions often describe long-run behavior quite well. Perfect rationality is frequently unnecessary for convergence to equilibrium; simple adaptive processes 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). Further, theoretical stability criteria are frequently employed to identify which equilibria are likely to emerge as the long-run outcome of adaptive dynamics (e.g., Hopkins and Seymour, 2002; Taylor and Jonker, 1978).

In strategic environments with a stable equilibrium, the long-run predictions of adaptive models often coincide with the predictions of equilibrium models, as adaptive models typically converge to stable equilibria. However, in strategic environments with unstable equilibria, adaptive dynamics may fail to converge, leading boundedly rational agents to exhibit persistently nonconvergent behavior and producing clean separation between the long-run predictions of adaptive models and those of equilibrium models. Characterizing long-run behavior in the absence of stable Nash equilibria is the primary focus of this paper. In particular the focus will be on identifying which aspects of long-run behavior are captured by equilibrium models, stability criterion, and adaptive dynamics.[1]

To test theoretical predictions from these three classes of models, this paper reports the results of experimental all-pay auctions[2] with highly unstable Nash equilibria. This class of games features a continuous range of pure strategies, so there a number of ways instability may occur. To experimentally test adaptive models of bidding behavior, these experiments provide subjects with continuous feedback, allow subjects to adjust their bids at will throughout each period, and have each subject earn mean-matching payoffs continuously over time. We examine two different all-pay auctions: one with two bidders and one prize

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

[2]Specifically we use all-pay auctions with n bidders who compete over $n-1$ equally valuable prizes. The top $n-1$ bidders each receive a prize, while the lowest bidder receives no prize. As in the standard all-pay auction, every bidder pays her bid.

(henceforth, "auction 1") and one with three bidders and two prizes (henceforth, "auction 2").

While auction 2 differs from auction 1 only in that it has one additional bidder and one additional prize, we show that it exhibits several key differences in theoretical predictions. For both auctions, Nash equilibrium involves bidders employing a mixed bidding strategy that leads to zero expected earnings. Because auction 2 has an additional prize, the equilibrium bid distribution has a higher mean in auction 2. The logit quantal response equilibrium (McKelvey and Palfrey, 1995) also predicts higher bids in auction 2. However, in contrast to the Nash prediction of zero expected earnings in both auctions, the logit quantal response equilibrium predicts positive payoffs for auction 2 and negative payoffs for auction 1. Each auction also exhibits very different stability properties. We show that auction 1 is globally neutrally stable (Theorem 3) and that auction 2 is highly unstable (Theorem 1). These results imply that adaptive dynamics exhibit convergence in auction 1, but exhibit non-convergent limit cycles in auction 2. We also show that these limit cycles are clockwise in the two-dimensional space of maximal earnings and optimal bids.

These theoretical results lead to several clear hypotheses for our experiment. We hypothesize that bids will be higher in auction 2 than auction 1 (Hypothesis 1), that payoffs will also be higher in auction 2 than auction 1 (Hypothesis 2), that auction 2 will exhibit greater instability than auction 1 (Hypothesis 3), and that non-convergent behavior will feature clockwise cycles (Hypothesis 4). Our results are remarkably consistent these hypotheses. We find higher average bids and payoffs in auction 2 than auction 1. Auction 2 exhibited significantly greater instability and significantly greater deviation from Nash equilibrium than auction 1. Bidding behavior in both auctions exhibited pronounced clockwise cycles identified by a positive cycle-rotation index.

However, observed bidding behavior also exhibits several significant departures from both Nash equilibrium and logit quantal response equilibrium. Average bids in both auction 1 and auction 2 were significantly lower than Nash equilibrium predictions. Consequently, average payoffs were significantly greater than zero in both auctions, contrasting with the logit quantal response equilibrium prediction of negative mean payoffs in auction 1.

As an additional treatment, orthogonal to our auction design, we also vary the way that information is presented to subjects. In the payoff information treatment subjects observe the instantaneous payoffs to all possible bids. In the social information treatment, subjects observe the current bids employed by all other subjects and their current instantaneous payoffs. This is merely a different way to present the same information as subjects can infer the same information about the current state of the game from either informational treatment. Our suspicion was that this differential presentation might alter bidding behavior, making subjects more likely to optimize under payoff information and making subjects more likely to imitate under social information. Since this is a purely behavioral intervention, both equilibrium and adaptive models predict no treatment effect from this differential presentation of information. (Hypothesis 5).

In auction 1, payoff information produces lower bids, higher payoffs and greater instability. In both auctions, the cycles predicted by adaptive models are more pronounced under payoff information. It is difficult to fully explain these results, but we do note two important points. The stronger effect of information on auction 1 rather than auction 2, may suggest global neutral stability is a weaker condition than positive definiteness. Additionally, the measure entirely determined by actual behavioral dynamics, the cycle-rotation index, is most affected by this treatment intervention. We thus see this treatment intervention an important vehicle for future research on adaptive dynamics.

From a big picture standpoint, these results tell us how to identify when equilibrium reliably characterizes aggregate behavior. When equilibria satisfy stability conditions, they describes long-run behavior well. When equilibria fail to satisfy stability conditions, they predict long run behavior poorly, but adaptive models still provide a reliable characterization long-run behavior.

The paper 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 Cason et al., we observe persistent, disequilibrium cycles that cannot be predicted by standard equilibrium-based models of game theory. Our paper extends these designs from games with two and three pure strategies to one with a continuum of pure strategies. Because players have a continuous range of pure strategies, there are several different ways that instability can occur. Our theory section identifies one particular cyclical pattern of disequilibrium behavior which is consistently predicted by adaptive models and our results section shows that this particular pattern is found within our data, providing stronger evidence for the predictive power of adaptive models.

The paper also contributes to the experimental literature on bidding behavior in all-pay auctions (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. To our knowledge, none of these papers conduct all-pay auctions in continuous time, a necessary condition for testing long-run predictions from adaptive dynamics.

This paper makes use of an already existing rich collection of theoretical models as the basis for many of its experimental tests, including theoretical predictions from both Nash equilibrium and the logit quantal response equilibrium described by Anderson et al. (1998), Baye et al. (1996), and McKelvey and Palfrey (1995). It also tests theoretical stability criteria described by Hopkins and Seymour (2002) and Taylor and Jonker (1978) and adaptive dynamics described by Fudenberg and Levine (1998). This paper also contributes the literature applying evolutionary stability concepts to experimental data (e.g., Crawford, 1991; Van Huyck et al., 1991). Our paper's novel experimental design provides a clean separation between the predictions of equilibrium models and those of adaptive dynamics in the all-pay auction, strongly rejecting the hypothesis of long-run behavioral stability. To our knowledge, these contributions are unique within the literature.

As a secondary treatment, our experiment also varies the way that information is presented in an attempt to learn more about imitative and optimization dynamics and their relation with stability. Previous literature has also investigated adaptive dynamics and stability without the use of continuous time. Much like our inconclusive finding on imitation, there appears to be little consensus on the predictive success or applicability of imitative models. Tests of imitative models using a wide variety of experimental designs have found evidence both in favor (Huck et al., 1999; Offerman et al., 2002; Feri et al., 2011) and against (Cheung and Friedman, 1998; Friedman et al., 2015) the predictive power of imitative over optimization models. Similarly, research on dynamic cognition and behavioral learning, which like our paper examines how individuals respond to information about games, finds 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. Perhaps closest to this informational treatment, Apesteguia et al. (2007) vary information in order to identify imitation. Unlike our design, their informational treatments vary whether players receive information about their direct competitors or other subjects with whom they do not directly compete.

This paper proceeds as follows: Section 2 describes the structure of each auction, presents their respective equilibrium predictions, and characterizes their respective adaptive dynamics. 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, bidders expend unrecoverable effort to compete over a limited number of prizes. Prizes are awarded to agents who expend the most effort, but every agent bears the cost of her effort, even if she does not win a prize. All-pay auctions are frequently employed as models for strategic environments that involve both conflict and unrecoverable 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).

Consider an all-pay auction where n bidders compete over $n-1$ prizes. Each bidder simultaneously selects a bid $b \in [0, w]$. All but the lowest bidder receives an identically valuable prize with publicly known value $v < w$. Every bidder must pay her bid. If multiple bidders are tied for the lowest bid, then the tie is broken at random. Let $L(s)$ denote the lowest bid under the strategy profile $s \in \mathbb{R}^n$, let $E(s)$ denote the number of bids equal to $L(s)$, and let $H(s) = (E(s) - 1)/E(s)$. Accordingly, the expected payoff to bidder i is given by:

$$\pi_i(s) = \begin{cases} v - b_i & \text{if } b_i > L(s) \\ vH(s) - b_i & \text{if } b_i \leq L(s) \end{cases}. \tag{1}$$

Hence the probability of receiving a prize by bidding $b \in [0, w]$ against opponents who employ the continuous mixed strategy F is given by

$$P(b|F) = \sum_{m=1}^{n-1} \binom{n-1}{m} F(b)^m (1 - F(b))^{n-m-1} \tag{2}$$

and the expected payoff to a bidder who bids $b \in [0, w]$ against opponents who employ the continuous mixed strategy F is given by

$$\pi(b|F) = vP(b|F) - b. \tag{3}$$

Further, the expected payoff to a bidder who employs the mixed strategy G against opponents who employ the mixed strategy F is given by

$$\pi(G|F) = \int_0^w \pi(b|F)\, dG(b) = v \int_0^w P(b|F)\, dG(b) - \int_0^w b\, dG(b). \tag{4}$$

In this paper, we focus on two special cases of this all-pay auction that exhibit very different stability properties. We refer to the all-pay auction with two bidders competing over one prize as "auction 1" and we refer to the all-pay auction with three bidders competing over two prizes as "auction 2." Hence the probability of winning auction 1 by bidding $b \in [0, w]$ against opponents who employ the mixed strategy F is given by

$$P(b|F) = F(b) \tag{5}$$

and the probability of winning auction 2 by bidding $b \in [0, w]$ against opponents who employ the mixed strategy F is given by

$$P(b|F) = 2F(b) - F(b)^2 \tag{6}$$

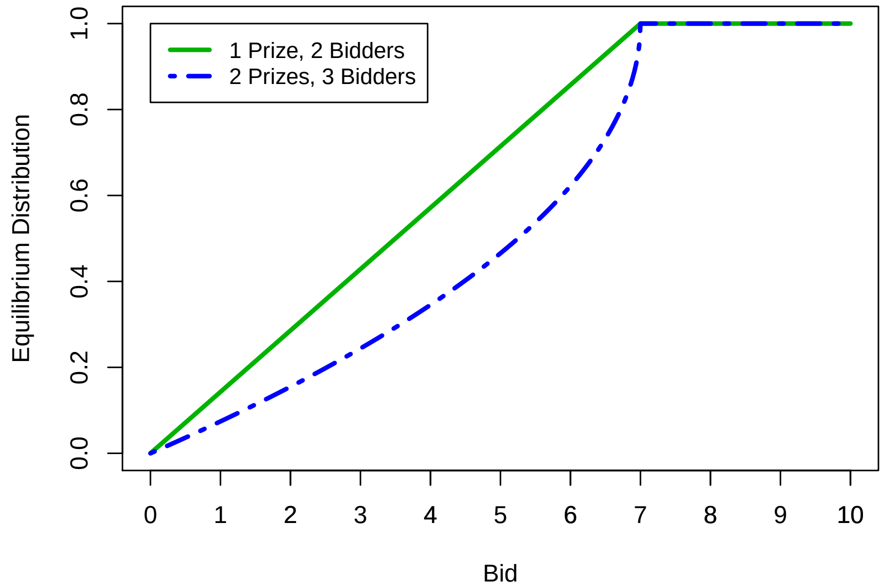

Figure 1: The Nash equilibrium bid distributions. The solid green line illustrates the equilibrium bid distribution for auction 1 and the dashed blue line illustrates the equilibrium bid distribution for auction 2.

2.1 Nash Equilibrium

Neither auction 1 nor auction 2 have a pure strategy Nash equilibrium, but they each have a unique symmetric mixed strategy Nash equilibrium.[3] The equilibrium bidding distribution for auction 1 was derived by Baye et al. (1996), and is given by

$$\Phi(b_i) = \frac{b_i}{v} \qquad \text{for } b_i \in [0, v] \tag{7}$$

The equilibrium bidding distribution for auction 2 was derived by Barut and Kovenock (1998)

$$\Phi(b_i) = 1 - \sqrt{1 - \frac{b_i}{v}} \qquad \text{for } b_i \in [0, v] \tag{8}$$

Figure 1 illustrates the Nash equilibrium bid distribution for each auction. The horizontal axis illustrates potential bids $b \in [0, w]$ and the vertical axis illustrates the equilibrium

[3] As discussed by Oprea et al. (2011), symmetric Nash equilibria tend to characterize the long-run outcome of repeated play under single population mean-matching more reliably than asymmetric Nash equilibria. Since our experimental design involves single population mean matching in continuous time, the symmetric Nash equilibrium is the most relevant theoretical prediction. That being said, investigating the emergence of asymmetric Nash equilibria in multi-population settings could provide an avenue for interesting future research.

cumulative probability $\Phi(b)$. Note that the Nash equilibrium bid distribution of auction 2 first-order stochastically dominates the Nash equilibrium bid distribution of auction 1, so Nash equilibrium predicts that bidders will bid more aggressively in auction 2.

2.2 Stability

To evaluate the stability of these Nash equilibria, we first consider the positive and negative definiteness conditions described by Hopkins and Seymour (2002). In a symmetric game with expected payoff function $\pi(b|F)$, an equilibrium mixed strategy distribution Φ with density function ϕ is said to be positive definite if the quadratic form

$$Q\left(Z\right)=\int D_{F}\pi(b|\Phi)Z(b)\,dZ(b) \tag{9}$$

is strictly positive for all $Z(b)=G(b)-\Phi(b)$ where $G\neq\Phi$ is an arbitrary non-equilibrium distribution with density function g. Here $D_{F}\pi(b|\Phi)$ denotes the linearization of the payoff function $\pi(b|F)$ at $F=\Phi$. Conversely, Φ is said to be negative definite if the quadratic form $Q(Z)$ is strictly negative. An equilibrium strategy that is neither positive definite nor negative definite is said to be indefinite. In all-pay auctions, this quadratic form can be written as

$$Q\left(Z\right)=\int_{0}^{w}\left.\frac{\partial\pi\left(b|F\right)}{\partial F\left(b\right)}\right|_{F=\Phi}Z\left(b\right)dZ\left(b\right) \tag{10}$$

Intuitively, an equilibrium strategy is negative definite if any sufficiently small deviation from equilibrium creates incentives that push behavior back towards equilibrium. Conversely, an equilibrium strategy is positive definite if any sufficiently small deviation from equilibrium creates incentives that push behavior farther away from equilibrium. To see why, suppose that bidders exhibit a small deviation from the equilibrium bid distribution Φ to a non-equilibrium bid distribution G. If G is sufficiently close to Φ then the expected payoff to a bid $b\in[0,w]$ against G is linearly approximated by

$$\tilde{\pi}\left(b|G\right)=\left.\frac{\partial\pi\left(b|F\right)}{\partial F\left(b\right)}\right|_{F=\Phi}Z\left(b\right) \tag{11}$$

and the quadratic form $Q(Z)$ can be written as

$$\int_0^w \tilde{\pi}\,(b|G)\,dZ(b) = \int_0^w \tilde{\pi}\,(b|G)\,dG(b) - \int_0^w \tilde{\pi}\,(b|G)\,d\Phi(b) = \tilde{\pi}(G|G) - \tilde{\pi}(\Phi|G). \tag{12}$$

Thus $Q(Z)$ provides a linear approximation for the difference between the expected payoff to the non-equilibrium strategy G against itself and the expected payoff to the equilibrium bidding strategy Φ against G. If $Q(Z)$ is strictly positive then the approximated payoff to the alternate bidding strategy G is strictly greater than the approximated payoff to the equilibrium bidding strategy. Conversely, if $Q(Z)$ is strictly negative then the approximated payoff to the alternate bidding strategy G is strictly less than the approximated payoff to the equilibrium bidding strategy. If the alternate strategy G is sufficiently close to the equilibrium strategy Φ then this local approximation accurately ranks the expected payoffs to each strategy.

Theorem 1. *The Nash equilibrium for auction 2 is positive definite.*

Proof. The Nash equilibrium bid distribution for auction 2 is given by $\Phi(b) = 1 - \sqrt{1 - b/v}$ for all $b \in [0, v]$ and the expected payoff function is given by $\pi(b|F) = 2vF(b) - vF(b)^2 - b$ so the linearization of the expected payoff function $\pi(b|F)$ at $F = \Phi$ is given by

$$\left.\frac{\partial \pi\,(b|F)}{\partial F\,(b)}\right|_{F=\Phi} = 2v(1 - \Phi(b)). \tag{13}$$

Now let $Z(b) = G(b) - \Phi(b)$ where $G \neq \Phi$ is an arbitrary non-equilibrium distribution with density function g. The quadratic form $Q(Z)$ from Equation (9) is given by

$$\begin{aligned} Q(Z) &= 2v \int_0^w \sqrt{1 - b/v}\; Z\,(b)\, dZ\,(b) \\ &= 2v \int_0^w \sqrt{1 - b/v}\; d\left(Z\,(b)^2/2\right) \\ &= 2v \left[\frac{1}{2} Z\,(b)^2 \sqrt{1 - b/v}\right]_0^w - v \int_0^w Z\,(b)^2 \; d\left(\sqrt{1 - b/v}\right) \\ &= -v \int_0^w Z\,(b)^2 \; \frac{d}{db}\left[\sqrt{1 - b/v}\right] db \end{aligned} \tag{14}$$

Thus is the quadratic form $Q(Z)$ strictly positive since $\frac{d}{db}\left[\sqrt{1 - b/v}\right] < 0$. □

Theorem 1 indicates that the unique Nash equilibrium of auction 2 is positive definite and in therefore unstable under a wide variety of adaptive models. Hopkins and Seymour (2002) show that every positive definite mixed strategy Nash equilibrium is an unstable point under all positive definite adaptive dynamics. Conversely, they show that every negative definite mixed strategy Nash equilibrium is a locally stable stationary point under all positive definite adaptive dynamics. Further, Hopkins (1999) shows that these stability results for positive definite adaptive dynamics extend to the both the best response dynamic and to any sufficiently precise perturbed best response dynamic.

Theorem 2. *The Nash equilibrium for auction 1 is indefinite.*

Proof. The Nash equilibrium bid distribution for auction 1 is given by $\Phi(b) = v/b$ for all $b \in [0, v]$ and the expected payoff function is given by $\pi(b|F) = vF(b) - b$ so

$$\left.\frac{\partial \pi\left(b|F\right)}{\partial F\left(b\right)}\right|_{F=\Phi} = v \tag{15}$$

Now let $Z(b) = G(b) - \Phi(b)$ where $G \neq \Phi$ is an arbitrary non-equilibrium distribution with density function g. Then the quadratic form $Q(Z)$ from Equation (10) can be written as

$$Q(Z) = v \int_0^w Z\left(b\right) dZ\left(b\right) = \frac{1}{2} v Z(w)^2 \tag{16}$$

Thus $Q(Z) = 0$ since we have $Z(w) = \Phi\left(w\right) - G\left(w\right) = 1 - 1 = 0$. □

Theorem 2 indicates that the Nash equilibrium for auction 1 is neither positive definite nor negative definite. It should be noted that positive and negative definiteness are local stability conditions that only describe the incentives created by small deviations from equilibrium. In contrast, a Nash equilibrium mixed strategy distribution Φ is said to be globally neutrally stable (Sandholm, 2010) if

$$\pi(\Phi|G) \geq \pi(G|G) \tag{17}$$

for any non-equilibrium mixed strategy distribution $G \neq \Phi$. Global neutral stability implies that the Nash equilibrium strategy does weakly better against any alternative strategy than this alternative strategy does against itself. So if a player's opponents were to employ the non-equilibrium mixed strategy G then she would be weakly better off employing the equilibrium strategy Φ than the non-equilibrium strategy G.

Theorem 3. *The Nash equilibrium for auction 1 is globally neutrally stable.*

Proof. The Nash equilibrium strategy for auction 1 is given by $\Phi(b) = v/b$ for all $b \in [0, v]$ and the expected payoff function is given by $\pi(b|F) = vF(b) - b$. Let $F \neq \Phi$ some arbitrary arbitrary non-equilibrium distribution on $[0, w]$. In this case, we have

$$\begin{aligned}
\int b \, dF(b) &= \int_{b=0}^{b=w} \int_{x=0}^{x=b} dx \, dF(b) \\
&= \int_{x=0}^{x=w} \int_{b=x}^{b=w} dF(b) \, dx \\
&= \int_0^w [1 - F(x)] \, dx
\end{aligned} \tag{18}$$

Now the expected payoff to the equilibrium strategy Φ against the non-equilibrium strategy F is given by

$$\begin{aligned}
\pi(\Phi|F) &= v \int F(b) \, d\Phi(b) - \int b \, d\Phi(b) \\
&= \int_0^v F(b) \, db - \frac{v}{2}
\end{aligned} \tag{19}$$

Conversely, the expected payoff to the non-equilibrium strategy F against itself is given by

$$\begin{aligned}
\pi(F|F) &= v \int F(b) \, dF(b) - \int b \, dF(b) \\
&= \frac{v}{2} - \int b \, dF(b) \qquad \text{since } F(X) \sim U[0,1] \text{ for } X \sim F(x) \\
&= \frac{v}{2} - \int_0^w [1 - F(b)] \, db \qquad \text{since } \int b \, dF(b) = \int_0^w [1 - F(b)] \, db \\
&\leq \frac{v}{2} - \int_0^v [1 - F(b)] \, db \\
&= \frac{v}{2} - v + \int_0^v F(b) \, db \\
&= \int_0^v F(b) \, db - \frac{v}{2} \\
&= \pi(\Phi|F)
\end{aligned} \tag{20}$$

Thus the equilibrium strategy Φ does weakly better against the non-equilibrium strategy F then the non-equilibrium strategy F does against itself. $\square$

Theorem 3 indicates that the Nash equilibrium strategy of auction 1 does at least as well against any non-equilibrium strategy then that non-equilibrium strategy does against itself. In contrast, Theorem 1 indicates that the Nash equilibrium strategy of auction 2 does worse against any sufficiently close alternative strategy then that alternative strategy does against itself. Together, these theorems suggest that the unique Nash equilibrium of auction 1 is fundamentally more stable than the unique Nash equilibrium of auction 2.

2.3 Adaptive Dynamics

The logit dynamic described by Fudenberg and Levine (1998) describes the evolution of a strategy distribution over time in a large population of agents. Agents in this population make asynchronous strategy adjustments where the timing of each agent's adjustments follows a homogeneous Poisson process. Unlike the perfectly rational agents described by Nash equilibrium, these agents do not always select a best responses but they are more likely to select strategies that yield higher payoffs. When the distribution of bids is given by the cumulative distribution function F then the likelihood that an adjusting agent will select the bid $b \in [0, w]$ is given by the logit response

$$\ell\left(b|F\right) = \frac{\exp \lambda \pi \left(b|F\right)}{\displaystyle\int_0^w \exp \lambda \pi \left(x|F\right) dx} \tag{21}$$

Hence the evolution of the bid distribution over time is governed by the differential equation $\dot{F}\left(b\right) = L\left(b|F\right) - F\left(b\right)$ where $L(b|F) = \int_0^b \ell(b|F)$. The parameter $\lambda \geq 0$ denotes the agent's sensitivity to payoff differences. When λ is large, agents have high precision and are sensitive to small differences in payoffs, so they are very likely to select strategies that yield high payoffs. As λ approaches infinity, agents become increasingly precise and the logit response approaches the best response $\arg\max_b \pi(b|F)$. When λ is small agents have low precision and are insensitive to small differences in payoffs, so they exhibit more randomness in their bidding behavior. When $\lambda = 0$, agents are completely insensitive to payoff differences and the logit response is uniformly distributed over the strategy space $[0, w]$.

The stationary points of the logit dynamic are the set of logit quantal response equilibria (McKelvey and Palfrey, 1995). As the precision parameter λ approaches infinity, the logit quantal response equilibrium approaches a Nash equilibrium. A closed form solution for the logit quantal response equilibrium of auction 1 is provided by Anderson et al. (1998) and is given by

$$L^*(b) = -\frac{1}{\lambda v}\log\left(1 - \frac{1-\exp(-\lambda b)}{1-\exp(-\lambda w)}(1-\exp(-\lambda v))\right) \tag{22}$$

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, we provide the logit quantal response equilibrium for auction 2 in Theorem 4.

Theorem 4. *The logit quantal response equilibrium bid distribution for auction 2 is given by* $L^*(b) = 1 - \frac{1}{\sqrt{\lambda v}}\text{erfi}^{-1}\left(\left[1 - \frac{1-\exp(-\lambda b)}{1-\exp(-\lambda w)}\right]\text{erfi}\left(\sqrt{\lambda v}\right)\right)$

Proof. In auction 2, the expected payoff to a bid $b \in [0, w]$ against the continuous mixed strategy distribution F is given by $\pi(b|F) = 2vF(b) - vF(b)^2 - b$, so the logit quantal response equilibrium bid distribution L must satisfy

$$\begin{aligned}
\frac{dL}{db} &= \frac{\exp\left(\lambda\left(2vL(b) - vL(b)^2 - b\right)\right)}{\int_0^w \exp\left(\lambda\left(2vL(y) - vL(y)^2 - y\right)\right)dy} \\
C\,dL &= \exp\left(\lambda\left(2vL - vL^2 - b\right)\right)db \\
C\exp\left(\lambda v\left(L^2 - 2L\right)\right)dL &= \exp(-\lambda b)\,db \\
C\exp(-\lambda v)\int \exp\left(\lambda v(1-L)^2\right)dL &= \int \exp(-\lambda b)\,db \\
\text{erfi}\left(\sqrt{\lambda v}(1-L)\right) &= C_1 - C_2\exp(-\lambda b) \\
L(b) &= 1 - \frac{1}{\sqrt{\lambda v}}\text{erfi}^{-1}\left(C_1 - C_2\exp(-\lambda b)\right) \qquad (23)
\end{aligned}$$

Where erfi denotes the imaginary error function and is given by $\text{erfi}(x) = \frac{2}{\sqrt{\pi}}\int_0^x \exp(x^2)dx$. Since bids are restricted to the closed interval $[0, w]$ we know that $L(0) = 0$ and $L(w) = 1$.

Solving these boundary conditions for the constants C_1 and C_2 obtains

$$
\begin{aligned}
C_1 &= \frac{\operatorname{erfi}\left(\sqrt{\lambda v}\right)}{\exp\left(-\lambda w\right) - 1} + \operatorname{erfi}\left(\sqrt{\lambda v}\right) \\
C_2 &= \frac{\operatorname{erfi}\left(\sqrt{\lambda v}\right)}{\exp\left(-\lambda w\right) - 1}
\end{aligned} \tag{24}
$$

Substituting the solutions for C_1 and C_2 into the logit quantal response equilibrium bid distribution for auction 2 yields

$$
L\left(b\right) = 1 - \frac{1}{\sqrt{\lambda v}} \operatorname{erfi}^{-1}\left(\left[1 - \frac{1 - \exp\left(-\lambda b\right)}{1 - \exp\left(-\lambda w\right)}\right] \operatorname{erfi}\left(\sqrt{\lambda v}\right)\right) \tag{25}
$$

□

Figure 2 illustrates the logit quantal response equilibrium mean bid under a variety of precision parameters. When agents are completely insensitive to payoff differences then bids are selected uniformly at random and the mean bid in both auctions is equal to $w/2$. As the precision parameter approaches infinity the logit quantal response equilibrium mean bid of each auction approaches the Nash equilibrium mean bid of each auction, $v/2$ in auction 1 and $2v/3$ in auction 2. Note that the logit quantal response equilibrium mean bid is larger in auction 2 than in auction 1 for every positive value of the precision parameter. Similarly, Figure 3 illustrates the logit quantal response equilibrium mean payoff under a variety of precision parameters. Note that the logit quantal response equilibrium mean payoff is always larger in auction 2 than in auction 1 for every possible value of the precision parameter.

Hopkins (1999) shows that if a unique mixed strategy Nash equilibrium is positive definite then the logit quantal response equilibrium is an unstable repeller of the logit dynamic under sufficiently high precision levels. Theorem 1 states that the unique mixed strategy Nash equilibrium of auction 2 is positive definite and therefore unstable while Theorem 3 states that the unique mixed strategy Nash equilibrium of auction 1 is globally neutrally stable. Thus sufficiently precise logit dynamics diverge to a persistent limit cycle in auction 2 but converge to the logit quantal response equilibrium in auction 1. Figure 4 depicts the projection of this limit cycle onto the two dimensional space with the most profitable bid on

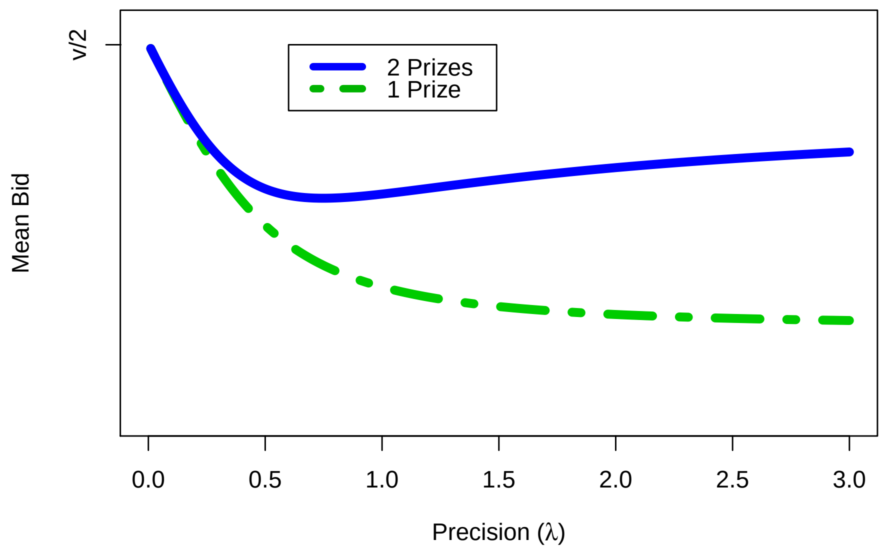

Figure 2: Mean bid under logit quantal response equilibrium.

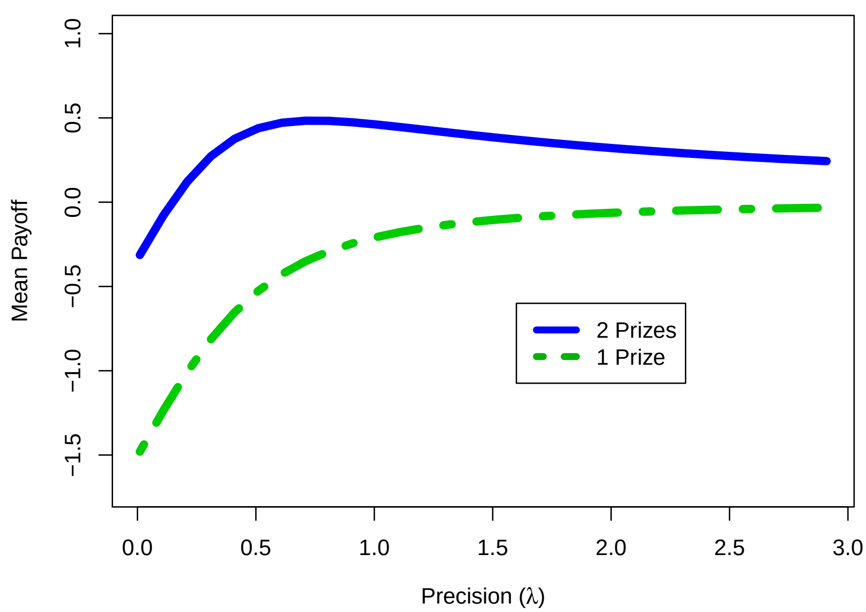

Figure 3: Mean payoff under logit quantal response equilibrium.

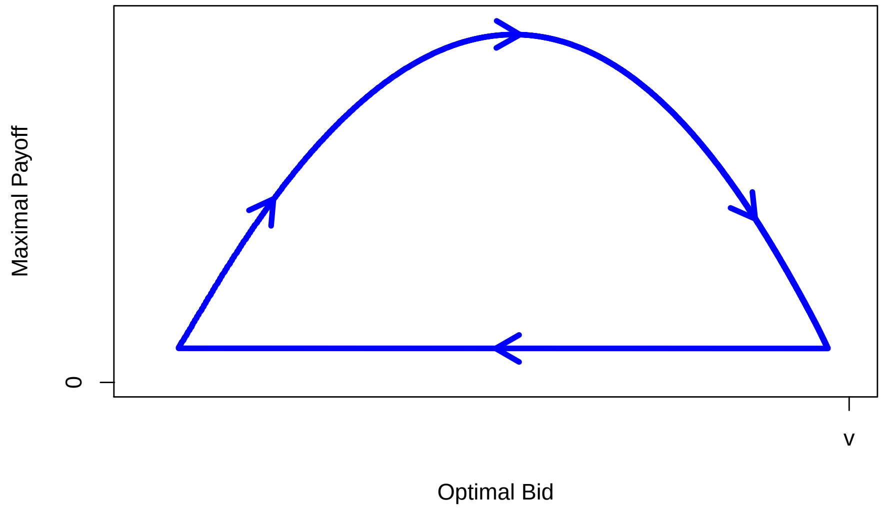

Figure 4: The limit cycle in auction 2 projected onto a two-dimensional space with the most profitable bid on the horizontal axis and the highest payoff on the vertical axis. The parameter values for this limit cycle are set to mirror the experimental setup for auction 2 and are given by $v = 7$, $w = 10$, and $\lambda = 10$.

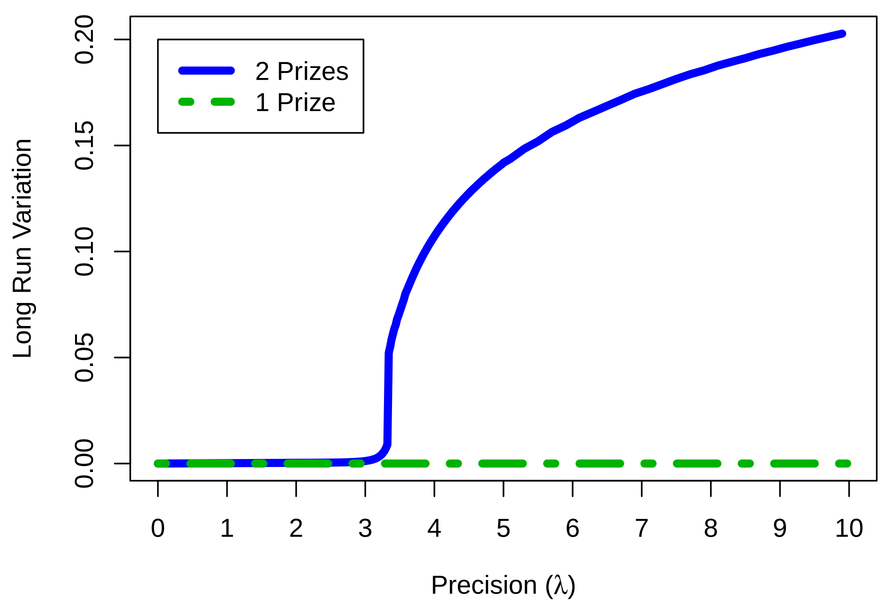

Figure 5: Long-run variation in the distribution of bids under the logit response dynamic.

the horizontal axis and the highest payoff on the vertical axis. Note that the figure depicts clockwise cycling.

Figure 5 depicts the long-run variation in the distribution of bids under a variety of precision parameters. Here variation is defined as the time average of the Chebyshev distance between the distribution of bids F_t in the population at time t and the long-run average bid distribution $\bar{F}(b) = \lim_{T\to\infty} \frac{1}{T}\int_0^T F_t(b)dt$. Formally, the long-run variation can be written as

$$\mathrm{Var}(F) = \lim_{T\to\infty} \frac{1}{T} \int_0^T \left\| \bar{F} - F_t \right\| dt, \tag{26}$$

where $\|G\| = \sup\{G(b) : b \in [0, w]\}$ denotes the uniform norm of the function G. Hence the long-run variation is equal to zero if the distribution of bids converges to a stationary point (auction 1, auction 2 under low λ) and the long-run variation is greater than zero if the distribution of bids converges to a limit cycle (auction 2 under high λ).

3 Experimental Design and Procedures

3.1 Design

This experiment implemented two all pay auctions, described in Section 2, with two bidders competing over one-prize (auction 1) and with three bidders competing over two prizes (auction 2). Subjects were endowed with $w = \$10$ and competed for prizes with value $v = \$7$. Subject bids were constrained to the interval $[0, w]$. Each session consisted of four 5 minute periods, during which 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. Using this interface subjects could select any bid amount in dollars and cents from \$0 to \$10.

To provide random rematching in continuous time, we employ a mean matching protocol (e.g. Cason et al., 2014; Oprea et al., 2011). A subject's instantaneous payoff was given by the expected value of her payoff from being randomly matched with other bidders in her session. Consistent with the theoretical framework described in Section 2, this means than each subject is competing against the distribution of bids currently employed by others in their session. This feature of mean matching is important for testing adaptive models because, without mean matching, collusion can easily occur in continuous-time games with relatively small group sizes (Friedman and Oprea, 2012; Friedman et al., 2015).

At the end of each session, subjects received the time average of their instantaneous payoff calculated in addition to a fifteen dollar show-up payment. 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. Subjects received continuous feedback regarding their instantaneous payoff throughout each period.

The experiment utilized a 2×2 experimental design. Two informational treatments were conducted with both the one prize auction and the two prize auction. Under the social-information treatment, each subject received real-time information regarding the bids and payoffs of every participant in her cohort. Under the payoff-information treatment, subjects observed her current payoff function, so she could directly observe the current payoffs to any possible bid. In both treatments, bids and payoffs were recorded at a rate of ten times per second.[4]

Figure 6 illustrates the experimental interface under the social-information and payoff-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 also displayed numerically at the bottom of the screen. In the social information treatment, the current bid and payoff of each other subject is represented by a red line. In the payoff information treatment, the subject's instantaneous payoff function is represented by a green line.

[4]Because payoffs are calculated ten times per second, one could argue this is actually a finitely repeated game. However, this approximation of continuous time is common in the literature (see Cason et al., 2014; Oprea et al., 2011). The continuous-time large population dynamics are arbitrarily well approximated by discrete-time finite population dynamics (Sandholm, 2010). Finally, it is unlikely people have the cognitive ability or physical reflexes to make adjustments as rapidly as this game is repeated, making the game effectively continuous from the subjects' standpoint.

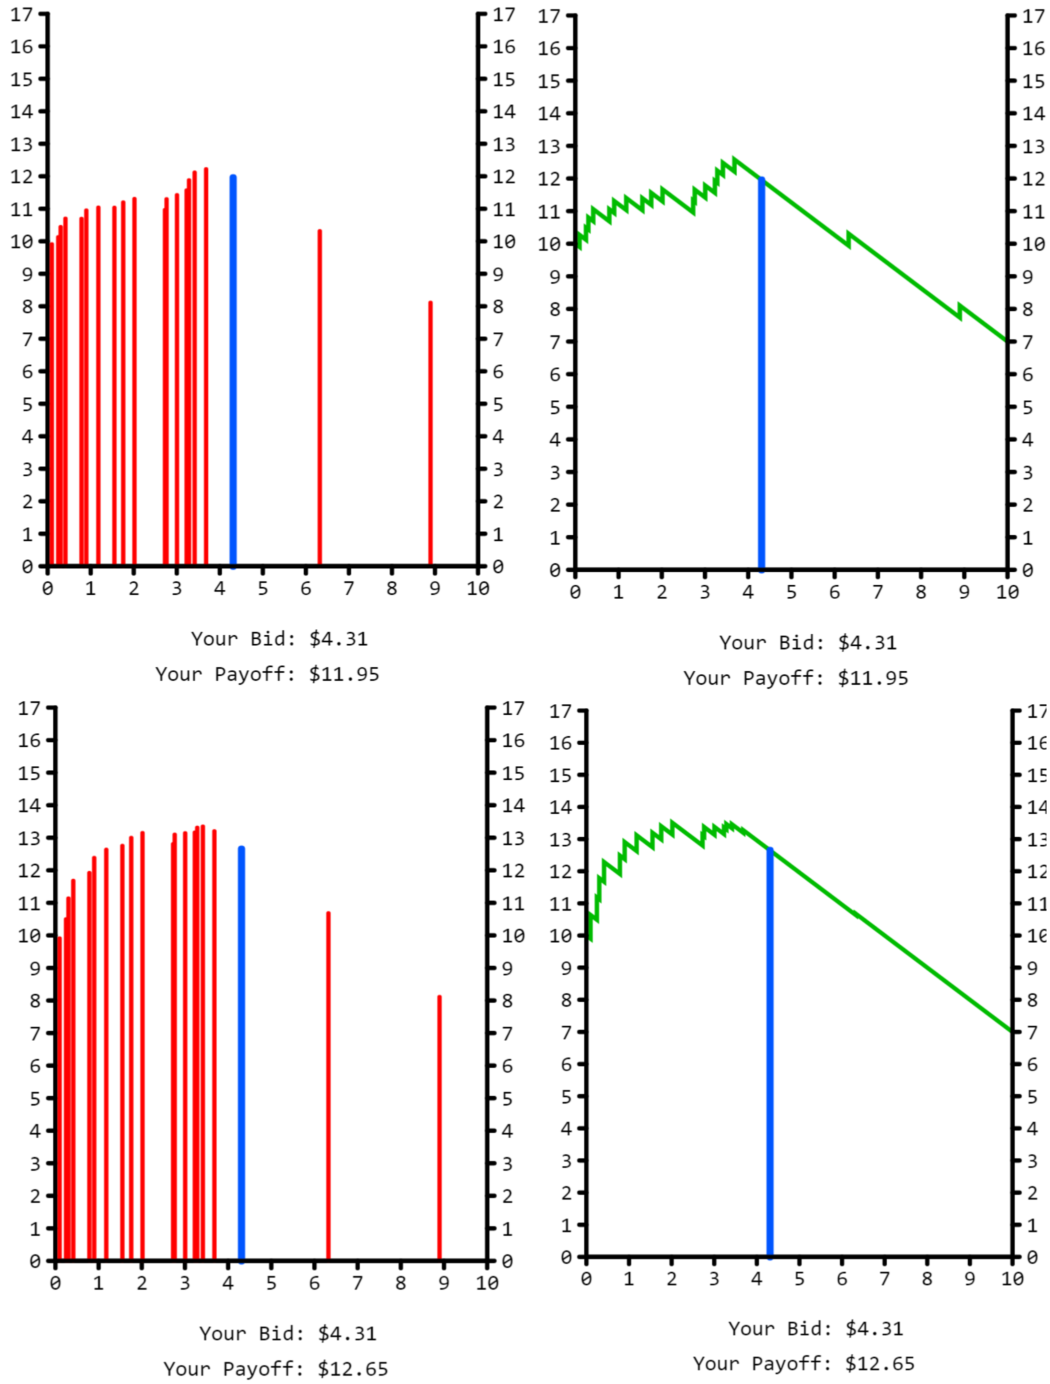

Figure 6: (a, top left) User interface under auction 1 with social information. (b, top right) user Interface under auction 1 with payoff information. (c, bottom left) User interface under auction 2 with social information. (d, bottom right) User interface under auction 2 with payoff information.

3.2 Procedures

Eight experimental sessions were conducted, two for each permutation of {payoff information, social information} × {auction 1, auction 2}. Each session was run with twenty experimental subjects. All 160 subjects were recruited from the Texas A&M undergraduate population using `econdollars.tamu.edu`, an ORSEE database (Greiner, 2015).

At the end of every session, each subject received the time average of their instantaneous payoff plus a fifteen dollar show-up payment. Subject earnings in auction 1 averaged $15.31 and $15.64 in the social and payoff information treatments, respectively. In auction 2, the earnings averaged $15.82 and $15.94 in the social and payoff information treatments, respectively. In equilibrium, average subject earnings would equal $15.00, so subjects received slightly above equilibrium earnings under both treatments. Sessions took place during October 2016. All sessions lasted less than one hour.

4 Hypotheses

As illustrated by Figure 1, the Nash equilibrium bid distribution of auction 2 first-order stochastically dominates the Nash equilibrium bid distribution of auction 1, so Nash equilibrium predicts that bidders will bid more aggressively in auction 2. Similarly, as illustrated by Figure 2, the mean bid for the logit quantal response equilibrium of auction 2 is higher than the mean bid for the logit quantal response equilibrium of auction 1 under any positive precision level. Accordingly, we have the following hypothesis.

Hypothesis 1. *Average bids will be higher in auction 2 than auction 1.*

Nash equilibrium predicts an expected payoff of zero under both auction 1 and auction 2. However, as illustrated in Figure 3, logit quantal response equilibrium predicts that the expected payoff to a bidder in auction 2 will be higher than the expected payoff to a bidder in auction 1, suggesting the following hypothesis.

Hypothesis 2. *Average payoffs will be higher in auction 2 than auction 1.*

Theorem 3 indicates that the Nash equilibrium strategy of auction 1 is globally neutrally stable, so it does at least as well against any alternative strategy then that alternative strategy does against itself. In contrast, Theorem 1 indicates that the the Nash equilibrium strategy of auction 2 is positive definite, so it does worse against any sufficiently close alternative strategy then that alternative strategy does against itself. Together, these theorems suggest that the Nash equilibrium of auction 2 is less stable than the Nash equilibrium of auction 1, resulting in the following hypothesis.

Hypothesis 3. *The empirical distribution of bids in auction 2 will exhibit greater instability than the empirical distribution of bids in auction 1.*

Figure 4 illustrates the clockwise limit cycles predicted by the logit dynamic in auction 2 when the distribution of bids is projected onto the two-dimensional space with the highest earning bid on the horizontal axis and the highest payoff on the vertical axis. Together with the local instability result from Theorem 1, the presence of these limit cycles results in the following hypothesis.

Hypothesis 4. *The empirical distribution of bids in auction 2 will exhibit clockwise cycles when projected onto the two-dimensional space with the highest earning bid on the horizontal axis and the highest payoff on the vertical axis.*

While the payoff and social information treatments feature different informational formats, one can deduce the payoff maximizing strategy from either display. Alternatively, one can infer and imitate the leading strategy from either display. Hence the aforementioned theoretical models yield identical predictions under both informational treatments, so we have the following hypothesis.

Hypothesis 5. *Subjects will exhibit no differences in behavior whether they observe payoff or social information.*

If Hypothesis 5 does not hold, a possible alternative is that providing payoff information will induces subjects to become optimizers and social information will induce subjects to become imitators. To the extent that these types of models predict differential behavior we may observe that differential behavior in the experiment. Alternatively, if subjects are inherently optimizers (or imitators) providing payoff (or social) information may cause them to perform more in line with theoretical predictions. However, we caution that because

Table 1: Summary statistics of bidding adjustments. Standard deviations are provided in parentheses.

	Overall	Auction 1 (1 prize, 2 bidders)		Auction 2 (2 prizes, 3 bidders)	
		Social Information	Payoff Information	Social Information	Payoff Information
Average Bid	3.688	3.378	3.125	4.086	3.865
	(1.795)	(1.820)	(1.708)	(1.700)	(1.793)
Dominated Bid	0.022	0.014	0.004	0.035	0.026
(bid above 7=1)	(0.146)	(0.117)	(0.064)	(0.183)	(0.159)
Minimum Bid	0.018	0.025	0.031	0.008	0.014
(bid below 0.25=1)	(0.132)	(0.157)	(0.172)	(0.091)	(0.116)
Absolute Change	0.894	0.824	1.033	0.774	0.944
in Adjustment	(1.083)	(1.046)	(1.084)	(1.005)	(1.146)
Payoff Change	0.413	0.173	0.423	0.385	0.556
with Adjustment	(0.881)	(0.561)	(0.750)	(0.880)	(1.044)
Observations	237,800	43,961	47,144	63,772	82,923
Sessions	8	2	2	2	2
Periods	32	8	8	8	8
Subjects	160	40	40	40	40
Time (in minutes)	160:00	40:00	40:00	40:00	40:00

there is no theoretical result behind these alternatives, they are speculative. Additionally, the wide variety of imitative and optimizing dynamics makes it difficult to isolate which type of dynamics are at work.

5 Results

Table 1 provides summary statistics for all subject bid adjustments over all sessions of the experiment. In total subjects made 237,8000 separate bid adjustments. That is, an individual subject made an average of 1486.25 bid adjustments over four five-minute periods, roughly one bid adjustment every 0.8 seconds. Bid adjustments were substantive, on average a subject adjusted their bid by \$0.89, roughly 9% of the strategy space. Subjects generally improved their payoff with each adjustment. On average, each adjustment increased instantaneous payoffs by \$0.17–0.56. Accordingly, subjects rarely adjusted to a dominated strategy (i.e., bidding over the value of the prize, 7) nor saw it necessary to avoid competing in the auction (i.e., bidding approximately 0).

We now turn to examining the theoretical predictions for the experiment. Both the

Table 2: Nash predictions and corresponding period-level empirical outcomes by auction. Standard errors are provided in parentheses. Stars indicate empirical difference between social and payoff information treatments is significant at 10%(*), 5% (**), and 1%(***) levels. Precision parameters are estimated by maximum likelihood, so the empirically estimated precision parameter for a given period is the precision parameter that maximizes the likelihood of the observed bid adjustments in that period under the logit dynamic.

	Nash Equilibrium Prediction		Empirically Observed Behavior	
	Auction 1 (1 prize, 2 bidders)	Auction 2 (2 prizes, 3 bidders)	Auction 1 (1 prize, 2 bidders)	Auction 2 (2 prizes, 3 bidders)
Mean Bid Amount	3.500	4.667	3.038 (0.067)	3.800*** (0.098)
Mean Payoffs	0	0	0.471 (0.066)	0.879*** (0.098)
Deviation from Time-Averaged Mean	0	0	0.193 (0.007)	0.283*** (0.009)
Deviation from Nash Equilibrium	0	0	0.232 (0.008)	0.389*** (0.016)
Cycle-Rotation Index	0	0	0.321 (0.036)	0.480*** (0.028)
Precision Parameter	∞	∞	1.106 (0.041)	0.972** (0.046)

Nash equilibrium and Logistic Quantal Response Equilibrium predict a higher mean bid for auction 2 than auction 1 (Hypothesis 1).

Result 1. *Mean bids are higher in auction 2 than auction 1. Payoffs are positive in both auctions, and higher in auction 2 than 1.*

Table 2 provides overall and by-treatment statistics on six key empirical values taken at the period level[5] and Nash Equilibrium predictions for comparison. As predicted by equilibrium models, the average bid in auction 1 (3.038) is lower than the average bid predicted in auction 2 (3.800). Both a t-test and a non-parametric Wilcoxon test find this difference to be statistically significant ($p < 0.001$). Both values are significantly lower than the Nash Equilibrium prediction ($p < 0.001$ for both t-test and Wilcoxon, both comparisons). Figure 7 provides a kernel density estimate of the bids in auction 1 and auction 2 along with

[5]For the remainder of the paper our analysis will be conducted at the period level, the unit we see as most appropriate to test our theoretical predictions.

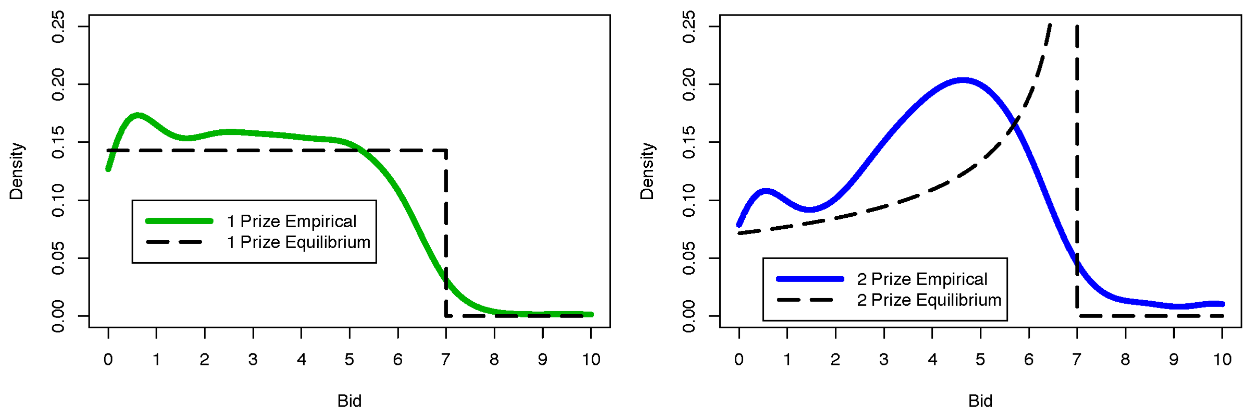

Figure 7: Bid Distribution of a 2 bidder, 1-prize all-pay auction (auction 1, left) and 3-bidder two prize all-pay auction (auction 2, right) compared to Nash Equilibrium predictions

the Nash Equilibrium predictions for each auction. Kolmogorov-Smirnov tests reveal that the empirical bid distributions are significantly different from each other and both are significantly different from their respective Nash Equilibrium bid distribution ($p < 0.001$, all comparisons).

Since the total number of prizes awarded in each auction is independent from bids, mean earnings are inversely related related to mean bids. Recall that the Nash Equilibrium predicts mean earnings of 0 in both auction 1 and auction 2. This prediction does not characterize the data well. Mean earnings in auction 1 (0.471) and auction 2 (0.879) are significantly positive and significantly greater in auction 2 ($p < 0.001$, parametric and non-parametric tests). The logit quantal response equilibrium prediction does slightly better; Figure 8 reveals that auction 2 is predicted to have positive earnings, but auction 1 is predicted to have negative earnings. Neither strategies nor payoffs in either auction match equilibrium predictions (though some comparative statics between the treatments are correct).

Moving beyond equilibrium predictions, both positive definiteness and global neutral stability criteria predict that auction 2 will exhibit greater instability than auction 1 (Hypothesis 3).

Result 2. *The empirical bid distribution in auction 2 exhibits both greater instability and greater deviation from Nash equilibrium than auction 1.*

Recall that each period provides 3000 discrete observations over time (5 minute periods

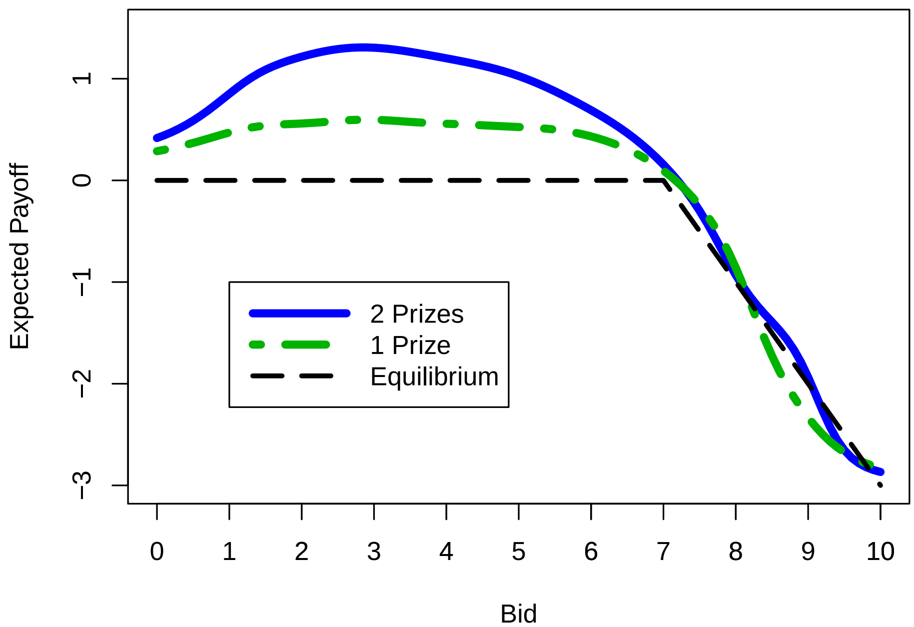

Figure 8: Predicted payoffs conditional on bid for a 2 bidder, 1-prize all-pay auction (auction 1) and 3-bidder two prize all-pay auction (auction 2) compared to nash equilibrium prediction.

$\times$ 0.1 second intervals). To measure variation in the empirical bid distribution over time, we first compute F_t, the empirical distribution of bids at at time time t. Next, we compute $\bar{F}$, the time-average distribution of bids over the entire period and we use the time-average of the Chebyshev distance between F_t and $\bar{F}$ as our measure of variation in the empirical distribution of bids. Alternatively, we compute F_A^*, the Nash equilibrium prediction for auction A and we use the time-average of the Chebyshev distance between F_t and $\bar{F}$ as our measure of deviation from equilibrium.

$$\begin{aligned} \text{Dev}(\bar{F}, F_t) &= \frac{1}{T}\sum_{t=0}^{T} \left\| F_t - \bar{F} \right\|, \text{ where} \\ \bar{F}(b) &= \frac{1}{T}\sum_{t=0}^{T} \| F_t(b) \| \\ \|G\| &= \sup\left\{ |G(x)| : x \in [0, w] \right\} \end{aligned}$$

The value $Dev(F, F_A^*)$ equivalently provides the deviation from Nash equilibrium in auction A.

Table 2 provides the means and standard errors for the aforementioned instability measure over all 32 periods in the experiment, separated by treatment. In accordance with Hypothesis Hypothesis 3, auction 2 exhibits significantly greater instability than auction 1. Consistent with theoretical predictions from stability criteria, (subsection 2.2) The instability measure for auction 2 (0.283) is significantly higher than for auction 1 (0.193) ($p < 0.001$, parametric and non-parametric tests), indicating that the empirical distribution of bids in auction 2 exhibits greater variation over time, a sign of instability. This result is consistent with theoretical results regarding the positive definiteness of auction 2 and the global neutral stability of auction 1.

While deviation from Nash equilibrium is not a measure of instability as the empirical distribution of bids might stabilize on a non-Nash distribution, it should be noted that both auctions exhibit significant deviation from Nash equilibrium, so neither auction matches the Nash predictions exactly. This finding is consistent with Result 1 which showed that the distribution on bids and corresponding average payoffs were significantly different than the Nash equilibrium value for both auctions. The deviations from equilibrium are also greater in auction 2 than 1 ($p < 0.001$, both parametric and non-parametric tests).

An interesting extension of this analysis is to observe the changes in the instability measure over time in each auction. Figure 9 displays the instability measure and the deviation from Nash measure over the four experimental periods in each session. Deviation from Nash equilibrium measure is relatively constant across both treatments. Consistent with Hypothesis 3, both measures of deviation have higher values for all periods of auction 2 than any period of auction 1.

Result 3. *Both auction 1 and 2 exhibit significant clockwise cycling. Cycling is more pronounced in auction 2.*

Adaptive dynamics predict cycling in auction 2 (Hypothesis 4), a prediction equilibrium models and stability criteria are not equipped to address. To measure cycles in the space bid distributions, we project this infinite dimensional space onto a more manageable two-dimensional space. To that end we employ a two-dimensional graph of maximum observed payoff and highest earning bid (see Figure 4). Replaying the data in real time on this graph

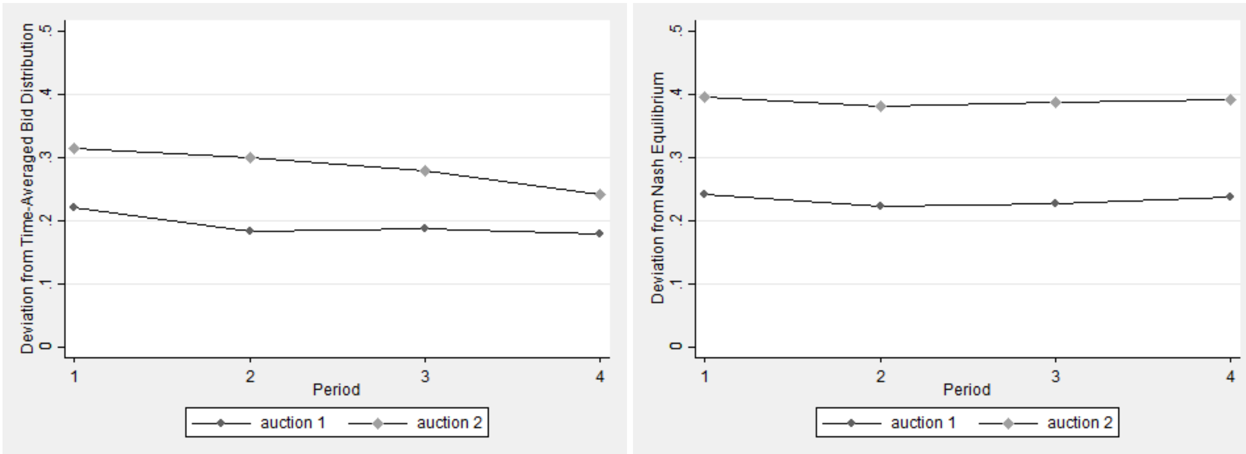

Figure 9: (a, left) Instability (i.e., deviation from time averaged bid distribution) by auction and period. (b, right) Deviation from Nash equilibrium prediction by auction and period.

we observe a great deal of persistent clockwise cycling. Figure 4 provides and example of such a cycle found between the 16 and 21 seconds of session 4.

To formally measure the strength of this clockwise cyclical behavior we take a Poincare section from the long-run average point in this space. This Poincare section is illustrated by the vertical line segment in Figure 10. Next we calculate the number of clockwise rotations and the number of counterclockwise rotations. Here a rotation occurs when the distribution crosses the Poincare section in the two-dimensional space. We then calculate the cycle rotation index described by (Cason et al., 2014)

$$CRI = \frac{\text{Clockwise Traversals} - \text{Counterclockwise Traversals}}{\text{Clockwise Traversals} + \text{Counterclockwise Traversals}}$$

If subjects exhibit exclusively clockwise rotations then cycle rotation index will equal 1. Conversely, if subjects exhibit exclusively counterclockwise rotations then cycle rotation index will equal -1. If subjects exhibit an equal number of clockwise and counterclockwise rotations the the cycle rotation index will equal 0. Similarly, if bids exhibit a stable distribution over time then the cycle rotation index will tend towards zero. If the cycle rotation index is significantly different from zero then the distribution of bids exhibits significant cyclical patterns. Adaptive dynamics predict that both auction 1 and auction 2 will exhibit significant clockwise cycling. However, in auction 1, these cycles are predicted to disappear in the

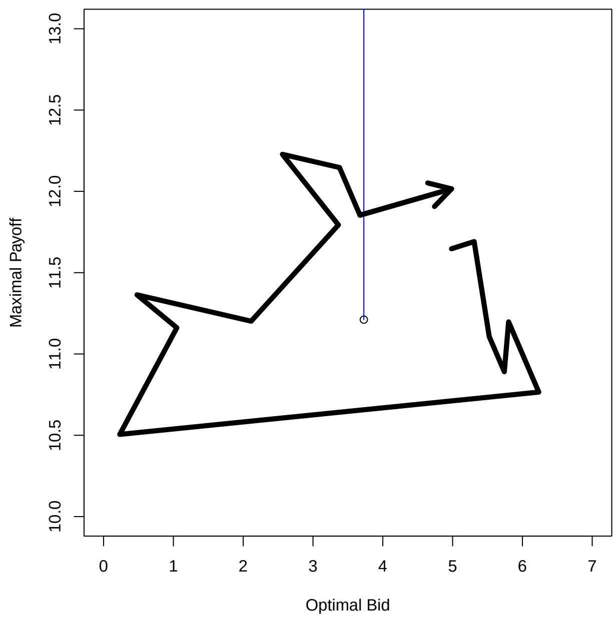

Figure 10: An example cycle on the two-dimensional graph of optimal bid and maximal payoff. Observed from seconds 16 to 21 of period 2, session 4.

limit as the bid distribution converges.

Table 2 provides the Cycle-Rotation Index for Auction 1 and Auction 2 by period. Both values are significantly different from 0 indicating pronounced clockwise cycling (parametric and non-parametric tests, $p < 0.001$). They are also significantly different from 1, the theoretical prediction of perfect clockwise cycling (parametric and non-parametric tests, $p < 0.001$), suggesting some noise in these cycles. Further, the cycle rotation index is significantly higher in auction 2, consistent with the theoretically predicted convergence of auction 1 under adaptive models.

Figure 11 reveals CRI averages over the four periods. For both auctions the CRI measure is significantly above 0 in all periods, indicating strong clockwise cycling. The CRI in auction 1 decreases in later periods while the CRI in auction 2 remains constant, consistent with the theoretical prediction of persistent limit cycles in auction 2 under adaptive models.

Result 4. *Showing payoff information rather than social information alters the empirically observed bid distribution in auction 1 but not auction 2. Only the cycle-rotation index is consistently affected by the presence of payoff information in both auctions.*

Similar to Table 2, Table 3 provides average and standard error values for several of our

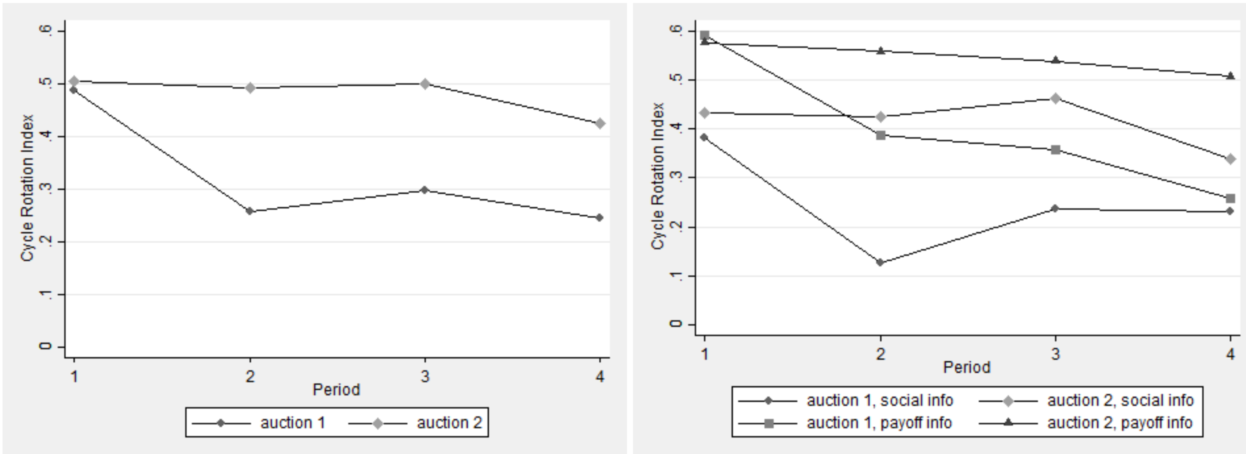

Figure 11: (a,left) Cycle-rotation index by auction and period. (b,right) Cycle-rotation index by auction, informational treatment, and period.

key values by auction type and informational treatment. We will now turn to address and compare these values now. Because our unit of observation is at the period level, and there are only two sessions of each period-treatment combination (8 observations total), we focus only on non-parametric hypothesis tests.[6]

In auction 1, payoff information is characterized by lower bids, higher payoffs, stronger cycling, and greater deviation from Nash equilibrium than social information. A permutations test finds each of these differences to be significant at the 5% level. A Kolmogorov-Smirnov test also finds the underlying bid distributions significantly different ($p < 0.001$).

Only one these informational treatment differences appear under auction 2. A permutation test does not reject the null hypothesis of equal bids, equal payoffs, or equal deviations from both Nash and the time average at the 10% level. The lone exception is the Cycle-Rotation index. Similar to auction 1, the Cycle-Rotation Index is higher under payoff information. Figure 11 shows the period specific Cycle-Rotation Index by auction and information treatment. Over all four periods within each auction, the payoff information treatment has a higher value on the CRI than its social information counterpart.

One final measure calculated is the Quantal Response Equilibrium Precision Parameter (λ). This measure was calculated by estimating a Logistic Quantal Response Equilibrium model for each period and treatment. A permutation test finds no difference between social

[6] Our table does provides standard errors for the inquisitive reader.

Table 3: Auction 1 and 2 period-level empirical outcomes by information type. Standard errors are provided in parenthesis. Stars indicate empirical difference between social and payoff information treatments is significant at 10%(*), 5% (**), and 1%(***) levels for a two-tailed non-parametric permutations test. Precision parameters are estimated by maximum likelihood, so the empirically estimated precision parameter for a given period is the precision parameter that maximizes the likelihood of the observed bid adjustments in that period under the logit dynamic.

	Auction 1 (1 prize, 2 bidders)		Auction 2 (2 prizes, 3 bidders)	
	Social Information	Payoff Information	Social Information	Payoff Information
Mean Bid Amount	3.203 (0.076)	2.874** (0.073)	3.741 (0.176)	3.858 (0.097)
Mean Payoffs	0.306 (0.076)	0.636** (0.073)	0.938 (0.176)	0.821 (0.098)
Deviation from Time-Averaged Mean	0.179 (0.008)	0.207** (0.008)	0.280 (0.011)	0.286 (0.015)
Deviation from Nash Equilibrium	0.205 (0.003)	0.259*** (0.006)	0.399 (0.024)	0.378 (0.021)
Cycle-Rotation Index	0.244 (0.037)	0.398** (0.050)	0.415 (0.030)	0.545** (0.035)
Precision Parameter	1.051 (0.065)	1.161 (0.047)	1.022 (0.058)	0.922 (0.070)

and payoff information treatment values in either auction 1 or 2 at the 10% level, suggesting that that subjects exhibit similar levels of precision under both informational treatments. Interestingly the estimated precision parameters differ slightly across auctions, bidders in auction 2 are estimated to have lower estimated precision than bidders in auction 1. We suspect that the higher precision of bidders in auction 1 may result from the greater stability of bid and payoffs in auction 1 relative to auction 2.

Simple regressions of these dependent variables on auction type, information treatment, their interaction, and a continuous variable for period number confirm these results (see Appendix Table A.1). Both the auction and information treatment coefficient are significantly positive in all regressions. The interaction term (auction $\times$ payoff information) coefficient is generally the opposite sign with similar magnitude to the information treatment variable. This relation suggests that there are differences between information treatments in auction 1, but not auction 2. An exception is the regression of the cycle-rotation index where the interaction term is essentially zero. The change in the CRI due to the information treatment can be interpreted as being the same across auctions.

6 Conclusions

This paper experimentally examines the differences between two all-pay auctions: one with two bidders and one prize (auction 1), the other with three bidders and two prizes (auction 2). We show that equilibrium models, stability criteria, and adaptive dynamics each predict large differences in bidding behavior between these two auctions. Nash equilibrium predicts that bidder in auction 1 will make lower bids. Logit quantal response equilibrium predicts that bidders in auction 1 will receive lower payoffs. Stability criteria predict greater stability in auction 1. Adaptive models predict convergence in auction 1 but persistent cyclical behavior in auction 2. To test these theoretical predictions and experimentally regarding long-run bidding behavior, we conducted continuous-time, mean-matching experiments with each of the aforementioned auctions.

Neither auction exhibits complete convergence to equilibrium in our data, but all four of the aforementioned theoretical predictions hold empirically. Most notably, auction 2 exhibited significantly greater instability, significantly greater deviation from equilibrium, and a significantly higher cycle-rotation index than auction 1, in line with theoretical predictions from adaptive models. Further, as predicted by logit quantal response equilibrium, average bids and payoffs were both significantly lower in auction 1 than auction 2.

The observed bidding behavior also exhibits several significant departures from both Nash equilibrium and the logit quantal response equilibrium. Average bids in both auction 1 and auction 2 were significantly lower than Nash equilibrium predictions. Consequently, average payoffs were significantly greater than zero in both auctions, contrasting with the logit quantal response equilibrium prediction of negative mean payoffs in auction 1. Further, while stability criteria correctly predict that auction 1 will be more stable than auction 2 (Theorem 1 and Theorem 3), they cannot characterize cyclical bidding dynamics as predicted by adaptive models.

The second, behaviorally based, treatment had more subtle results. In theory, there should be no difference between supplying subjects with payoff or social information in this game as each information treatment provides subjects with the same information in a different format. Further, the best response, logit response, and any imitative response could be calculated under either informational treatment. However, we observe some key differences in variables of interest across informational treatments. In auction 1, payoff information is associated with significantly lower bids, higher payoffs and greater instability. None of these differences are found to be significant in auction 2. Across both auctions, the cycle rotation index is significantly higher under payoff information, suggesting that the payoff information treatment is associated with more robust clockwise cycles. It also suggests that payoff information brings subjects closer to the predictions from adaptive dynamics which yield a CRI value of 1. Surprisingly, the logit precision parameter (λ), a measure of payoff responsiveness, is not significantly different across either type of information treatment.

The intent of using payoff (social) information was to make optimization (imitative) easier for subjects and induce those dynamics in subject play. It is unclear whether this occurred.[7] However, there is suggestive evidence that the differential presentation of information altered dynamics in some ways. We do note three interesting relations between theory and results and speculate on their meaning.

First, the stability condition of auction 1, global neutral stability, is different than the instability condition of auction 2, positive definiteness. Since auction 2 appears more robust to the information treatment, this may indicate that the latter theoretical property is stronger. Second, the cycle-rotation index is the one measure sensitive to the information treatment in both auctions. It is also most related to models of adaptive dynamics. Perhaps this means the information treatments are effective at altering adaptive dynamics but not always effective at altering other equilibrium-related measures. Finally, the cycle-rotation index is closer to 1 in both auctions under the payoff information treatment. This might point to optimization rather than imitation being representative of the way individuals process and respond to information. However, given the similarity in subject performance across

[7] One approach might be to simulate imitative and optimization models under the conditions of auction 1 and auction 2 to see whether the empirically observed differences occur. This approach is problematic for several reasons. There are many different imitative and optimization models that we could use. Slight changes in assumptions and parameters could lead to various empirical differences on our variables of interest. Second, our initial investigation into these methods led to predictions contrary to our data. For instance, simulations generally predict greater instability in the distribution of bids under imitative dynamics than under optimization dynamics, contrary to what is observed across our social and payoff treatments.

information treatments under a whole host of other measures, it will require future research to arrive at any definitive conclusion on this matter.

Regardless, our main results concerning stability and cycling help us understand when we can use Nash equilibrium to reliably predict aggregate behavior in games. If equilibrium satisfies stability conditions, equilibrium models provide a good estimate of long-run aggregate behavior. However, if equilibrium models fail to satisfy stability conditions, they are unreliable but adaptive models still provide a reliable characterization long-run behavior. We invite further research on this topic.

References

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.

Apesteguia, J., Huck, S., and Oechssler, J. (2007). Imitation theory and experimental evidence. *Journal of Economic Theory*, 136(1):217–235.

Barut, Y. and Kovenock, D. (1998). The symmetric multiple prize all-pay auction with complete information. *European Journal of Political Economy*, 14(4):627–644.

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.

Baye, M. R., Kovenock, D., and De Vries, C. G. (1996). The all-pay auction with complete information. *Economic Theory*, 8(2):291–305.

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

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.

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

Crawford, V. P. (1991). An evolutionary interpretation of van huyck, battalio, and beil's experimental results on coordination. *Games and Economic behavior*, 3(1):25–59.

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.

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.

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.

Friedman, D. and Oprea, R. (2012). A continuous dilemma. *The American Economic Review*, 102(1):337–363.

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

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

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.

Hopkins, E. and Seymour, R. M. (2002). The stability of price dispersion under seller and consumer learning. *International Economic Review*, 43(4):1157–1190.

Huck, S., Normann, H.-T., and Oechssler, J. (1999). Learning in cournot oligopoly–an experiment. *The Economic Journal*, 109(454):80–95.

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.

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

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.

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

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

Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *The Quarterly Journal of Economics*, 106(3):885–910.

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

A Additional Tables and Figures

Table A.1: Linear regressions of key dependent variables on treatments and interactions. All standard errors are heteroskedasticity-robust.

	Mean Bid	Mean Payoff	Deviation From Time-Averaged Mean	Deviation From Nash Equilibrium	Cycle-Rotation Index	Precision Parameter λ
auction 2	0.538***	0.631***	0.102***	0.619***	0.193***	-0.028
	(0.163)	(0.164)	(0.022)	(0.046)	(0.025)	(0.075)
payoff information	-0.329***	0.330***	0.028**	1.587***	0.053***	0.110
	(0.091)	(0.091)	(0.010)	(0.325)	(0.007)	(0.067)
auction 2 × payoff information	0.446***	-0.446**	-0.022	-0.400	-0.074	-0.210*
	(0.209)	(0.210)	(0.015)	(0.358)	(0.033)	(0.116)
period	-0.115**	0.114*	-0.018***	-0.001	-0.046***	0.051
	(0.058)	(0.058)	(0.004)	(0.009)	(0.016)	(0.030)
observations	32	32	32	32	32	32
r-squared	0.706	0.497	0.867	0.778	0.648	0.315