Playing the Field in All-Pay Auctions*

Daniel G. Stephenson†
Department of Economics
Virginia Commonwealth University

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

June 15, 2020

Abstract

We provide the first examination of all-pay auctions using continuous-time protocols, allowing subjects to adjust their bid at will, observe payoffs almost instantaneously, and gain more experience through repeated-play than in previous, discrete-time, implementations. Unlike our predecessors—who generally find overbidding—we observe underbidding relative to Nash equilibrium. To test the predictions of evolutionary models, we vary the number of bidders and prizes across treatments. If two bidders compete for a single prize, evolutionary models predict convergence to equilibrium. If three bidders compete for two prizes, evolutionary models predict non-convergent cyclical behavior. Consistent with evolutionary predictions, we observe cyclical behavior in both auctions and greater instability in two-prize auctions. These results suggest that evolutionary models can provide practitioners in the field with additional information about long-run aggregate behavior that is absent from conventional models.

*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. This paper benefitted from helpful comments by Tim Cason, Catherine C. Eckel, Daniel Fragiadakis, Dan Friedman, Dan Kovenock, Ryan Oprea, as well as the continual support and feedback the experimental research team at Texas A&M. We are also grateful for the insights given by attendees of the 2013 North American and World Meetings of the Economic Science Association.

†Corresponding author. Please contact at stephensod@vcu.edu.

‡Email: alexbrown@tamu.edu.

1 Introduction

Nash equilibrium often provides an approximate characterization of long-run behavior, but the short-run behavior of inexperienced agents often bears little resemblance to equilibrium predictions (see Davis and Holt, 1993). In all-pay auctions, Nash equilibrium can involve non-uniform mixed strategies over a continuum of actions. Considering the complexity of such equilibria, it seems unsurprising that inexperienced subjects would fail to play it. Previous studies find that subjects consistently overbid relative to Nash equilibrium and often play dominated strategies by overbidding beyond the actual value of the prize (Dechenaux et al., 2014; Gneezy and Smorodinsky, 2006; Lugovskyy et al., 2010). After many periods of experience, overbidding often decreases, but is not eliminated (Davis and Reilly, 1998; Lugovskyy et al., 2010). It remains unclear how much overbidding would persist in the long run as subjects continue to gain experience.[1] Experimental procedures designed to match the assumptions of long-run play, such as continuous-time experimental protocols, may help address this question.

This paper reports the first experimental investigation of the all-pay auction in continuous time. Subjects could adjust their bids at will, receive almost instantaneous feedback, and earn mean-matching payoffs continuously over time. Continuous-time experimental protocols effectively reduce the length of discrete periods to near instants. Mean-matching protocols provide subjects with information about their expected payoff from being randomly matched against others. Together, these experimental protocols provide subjects with far more experience in a fraction of the time relative their conventional discrete-time counterparts, allowing boundedly rational subjects to more easily assess the expected profitability of their strategies.

In contrast to previous experimental studies of the all-pay auction—all of which used discrete-time protocols—mean bids *did not* exceed the Nash predictions. On average, subjects slightly underbid relative to Nash and earned positive payoffs. Subjects rarely played dominated strategies; bids only exceeded the value of the item in roughly 3% of observations.

[1] Inexperience is one of many explanations for overbidding relative to Nash equilibrium in all-pay auctions (see Dechenaux et al., 2014). Preference-based explanations (e.g., joy of winning an item) may cause overbidding to continue in the long-run. We explore this issue in more detail in our concluding section.

These initial findings suggest that previous overbidding results may not apply to settings where individuals accumulate a large amount of experience, a result that may have bearing in the field. As Dechenaux et al. (2014) notes, empirically studying applications of the all-pay auction is often infeasible as the cost of effort is frequently unobservable. While experiments are imperfect substitutes for field settings, they provide an alternative method of investigation. If competitors in the field have extensive experience, we should cautiously interpret the overbidding observed in past discrete-time experimental studies where subjects received far less experience. The use of continuous time experimental protocols may provide a bridge from these conventional discrete-time laboratory settings to field settings where agents accumulate extensive experience.

While average bids did not exceed Nash predictions, aggregate bidding behavior did exhibit strong cyclical patterns that run contrary to Nash predictions. Evolutionary models predict cyclical behavior in all-pay auctions because outbidding one's opponent is optimal only if their bid is below the value of the prize. Once bids are at or above the value of the prize, the best response is to bid zero and the process repeats. This results in a cyclical pattern of gradual bid increases followed by sharp decreases. Consequently, all-pay auctions provide a particularly informative test of evolutionary models as their evolutionary predictions can be very different from their equilibrium predictions. The new application of continuous-time and mean-matching experimental protocols to this environment mirrors the structure of evolutionary models. We test evolutionary dynamics in two distinct auctions as between-subjects experimental treatments. In auction 1, two bidders compete for a single prize. In auction 2, three bidders compete for two prizes. Nash equilibrium predicts a fixed distribution of bids in each auction. In contrast, evolutionary models predict cyclical behavior in both auctions, convergence to a fixed point in auction 1, and persistent non-convergence in auction 2.[2]

Aggregate behavior exhibited persistent cycles in both auctions and greater instability in auction 2 than 1. Because these results are predicted by evolutionary models but not

[2]Time-averaged evolutionary predictions are similar to the Nash predictions, but dynamic evolutionary predictions are very different from the Nash predictions.

by conventional fixed-point models,[3] they illustrate how evolutionary models can provide additional substantive information not found in conventional models. This contribution is important for two distinct reasons. First, the all-pay auction is applicable to a rich set of concrete field environments 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). For practitioners who encounter such settings in the field, evolutionary models can identify which strategic environments are likely to exhibit long-run behavior that approaches equilibrium predictions. Second, while other studies observe cycles and instability in games with a relatively small number of actions (i.e., 2 or 3), the continuous-action space of the all-pay auction allows us to provide evidence for the applicability of evolutionary models to a much wider class of environments (see section 2 for more detail). This includes a variety of concrete settings where evolutionary models are employed such as finance (Hens and Schenk-Hoppé, 2005), bargaining (Abreu and Sethi, 2003), and industrial organization (Alós-Ferrer et al., 2000). When evolutionary models predict convergence, our results suggest that equilibrium models can often provide a useful characterization of long-run aggregate behavior. Conversely, when evolutionary models predict non-convergence, our results suggest that equilibrium models often fail to characterize long-run aggregate behavior. Even when equilibrium models are unreliable, our results suggest that evolutionary models can provide practitioners in the field with a useful characterization of behavioral dynamics.[4]

The remainder of this paper proceeds as follows: Section 2 discusses the related literature. Section 3 presents the theoretical predictions. Section 4 describes the experimental design

[3]Fixed-point models, including both Nash equilibrium and quantal response equilibrium (McKelvey and Palfrey, 1995), identify invariant points of an operator on the strategy space. Nash equilibria are fixed points of the best response and quantal response equilibria are fixed points of the quantal response. Consequently, fixed-point models never predict instability in the distribution of strategies. Equilibrium selection models can address relative stability across multiple equilibria, but are less useful in settings with a unique equilibrium. In contrast, evolutionary models describe a dynamic adjustment process that can explain persistently non-convergent cyclical patterns in the distribution of strategies across a population.

[4]For example, consider the design of a contest structure to allocate grant funding. Several distinct contest structures may yield identical equilibrium predictions. Evolutionary models can help policymakers identify which of these structures is most likely to induce convergence on the desired equilibrium. Conversely, a policymaker who fails to consider evolutionary models may select a contest structure that has desirable equilibrium properties but induces undesirable non-convergent behavior. We elaborate on this idea in our concluding section.

and procedures. Section 5 presents the experimental hypotheses. Section 6 provides the main results. Section 7 discusses our main findings and provides context.

2 Related Literature

This paper reports the first experimental investigation of the all-pay auction under continuous-time protocols. As such, it connects two strands of literature: experimental investigations of the all-pay auction and experimental studies of continuous-time dynamics.[5] Previous experimental studies of the all-pay auction (e.g., Gneezy and Smorodinsky, 2006; Lugovskyy et al., 2010; Ernst and Thöni, 2013) conduct a sequence of discrete periods in which subjects secretly select their bids and the single highest bidder receives a price. With the notable exception of Potters et al. (1998), these experiments are generally characterized by persistent overbidding relative to the Nash equilibrium. We are the first to conduct such experiments in continuous time and do not observe a similar result.

The paper also contributes to the experimental literature on continuous-time evolutionary dynamics (i.e., Benndorf et al., 2016; Bigoni et al., 2015; Cason et al., 2014, 2020; Oprea et al., 2011; Stephenson, 2019). Consistent with Cason et al. (2014) and Stephenson (2019), we find persistent cyclical behavior when adaptive models are unstable. Unlike previous studies, this paper identifies cyclical behavior in games with a continuum of pure strategies, instead of simple bimatrix games with two or three strategies. This distinction necessitates an important procedural difference. Cason et al. (2014) and Stephenson (2019) allow subjects to directly select mixed strategies. Because visualizing the 1000 dimensional space of probability distributions over these actions is infeasible,[6] our experiment allowed subjects to directly select bids. As in evolutionary models, mixed strategies occur in our experiment only as distributions of pure strategies over a population of subjects. Evolutionary predictions largely hold in our experiment, providing a precedent for further testing of evolutionary models in games with mixed-strategy equilibria over large numbers of actions.

Finally, this paper makes extensive use of the theoretical literature. Both Nash equilibrium

[5]See Dechenaux et al. (2014) and Brown and Stephenson (2019) for a survey on each topic, respectively.
[6]The all-pay auction has a continuous strategy space which we discretize into 1001 actions (see section 4).

and logit quantal response equilibrium (Anderson et al., 1998; Baye et al., 1996; McKelvey and Palfrey, 1995) provide a basis for experimental tests of fixed-point models. We also test the stability criteria provided by Hopkins and Seymour (2002) and Taylor and Jonker (1978) and the logit dynamics described by Fudenberg and Levine (1998). Our experimental design provides a clean separation between the predictions of equilibrium models and those of evolutionary dynamics in the all-pay auction, strongly rejecting the hypothesis of aggregate behavioral stability in the all-pay auction with two prizes.

3 Theory

Consider an all-pay auction where n bidders compete over $n-1$ indivisible prizes. Each bidder i simultaneously selects a bid $s_i \in [0, w]$. All but the lowest bidder receives an identically valuable prize with publicly known value $v \leq w$. If multiple bidders are tied for the lowest bid, then the tie is broken at random. Let $L(s)$ denote the lowest bid in the pure 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:

$$\pi\left(s_i | s_{-i}\right) = \begin{cases} v - s_i & \text{if } s_i > L(s) \\ vH(s) - s_i & \text{if } s_i \leq L(s) \end{cases}. \tag{1}$$

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

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

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

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

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

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

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

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

3.1 Fixed-Point Models

Fixed-point models characterize behavior by invariant points of operators on the strategy space. Nash equilibria are fixed points of the best response. Logit quantal response equilibria are fixed points of the logit quantal response. In general, such fixed-points may be pure strategy profiles or non-trivial mixed strategy profiles, but they must be a single strategy profile.

3.1.1 Nash Equilibrium

Neither auction 1 nor auction 2 has a pure strategy Nash equilibrium; they each have a unique Nash equilibrium in symmetric mixed strategies (Barut and Kovenock, 1998). The equilibrium bidding distribution for auction 1 was derived by Baye et al. (1996), and is given by

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

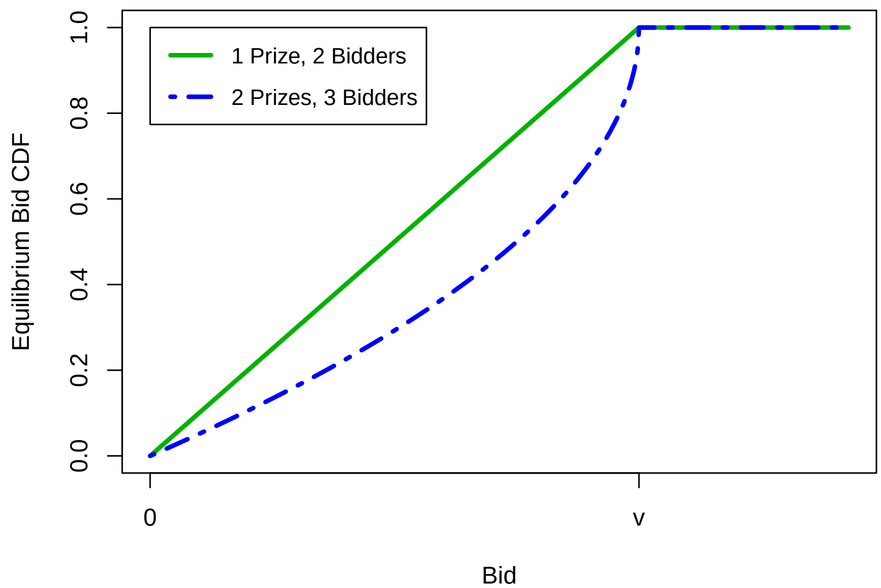

Figure 1: The Nash equilibrium bid distributions. The solid green line illustrates the equilibrium bid distribution for auction 1 and the dashed blue line illustrates the equilibrium bid distribution for auction 2.

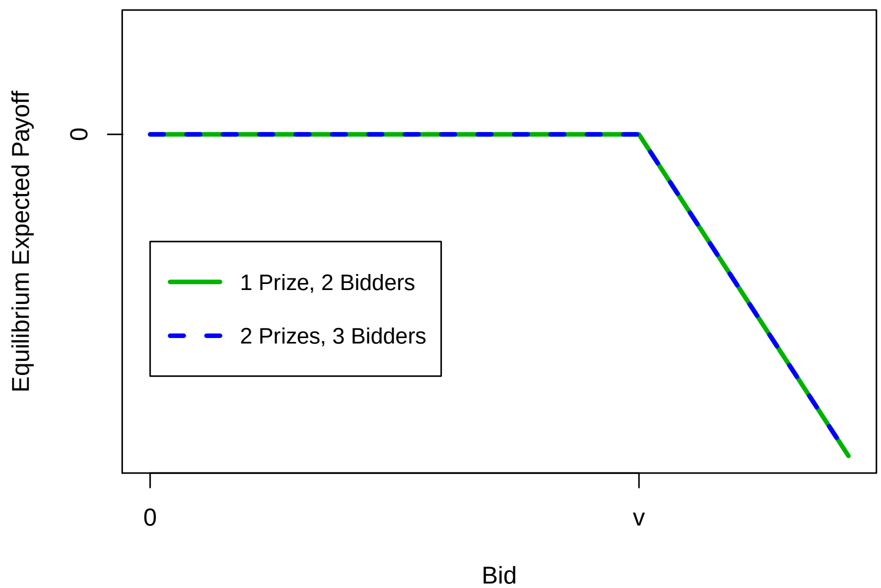

Figure 2: The Nash equilibrium expected payoff functions. The solid green line illustrates the equilibrium expected payoff function for auction 1 and the dashed blue line illustrates the equilibrium expected payoff function for auction 2.

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

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

Figures 1 and 2 illustrate the Nash equilibrium bid distributions and expected payoff functions respectively. The horizontal axis illustrates potential bids $b \in [0, w]$ and the vertical axes illustrates the equilibrium cumulative probability and the equilibrium expected payoff respectively. Equilibrium expected payoffs are zero[7] in both auctions, but the Nash equilibrium bid distribution of auction 2 first-order stochastically dominates[8] the Nash equilibrium bid distribution of auction 1, implying that bidders will bid more aggressively in auction 2

3.1.2 Logit Quantal Response Equilibrium

Unlike the perfectly rational agents described by Nash equilibrium, the agents described by logit quantal response equilibrium do not always select a best response, 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

$$\ell(b|F) = \frac{\exp \lambda \pi (b|F)}{\int_0^w \exp \lambda \pi (x|F)\, dx} \tag{9}$$

A logit quantal response equilibrium is a fixed-point L^* of the logit response such that

$$L^*(b) = \mathcal{L}(b|L^*) = \int_0^b \ell(x|L^*)\, dx \tag{10}$$

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. When λ is small,

[7]Since the expected payoff from bidding zero is zero, the expected payoff from other bids in the support of the Nash equilibrium must also be zero.

[8]Since $\sqrt{x} > x$ for all $x \in (0, 1)$, we have $\frac{b_i}{v} > 1 - \sqrt{1 - \frac{b_i}{v}}$ for all $b_i \in (0, v)$.

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 response are the 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

$$L^* (b) = -\frac{1}{\lambda v} \log \left(1 - \frac{1 - \exp(-\lambda b)}{1 - \exp(-\lambda w)} (1 - \exp(-\lambda v)) \right) \tag{11}$$

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 proposition 1.

Proposition 1. *The logit quantal response equilibrium bid distribution for auction 2 is*

$$L^* (b) = 1 - \frac{1}{\sqrt{\lambda v}} \operatorname{erfi}^{-1} \left(\left[1 - \frac{1 - \exp(-\lambda b)}{1 - \exp(-\lambda w)} \right] \operatorname{erfi} \left(\sqrt{\lambda v} \right) \right) \tag{12}$$

Proof. See appendix on page 34. □

Figure 3 illustrates the logit quantal response equilibrium mean bid under a variety of precision parameters. If 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. 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 4 illustrates the logit quantal response equilibrium mean payoff under a variety of precision parameters. The logit quantal response equilibrium mean payoff is larger in auction 2 than in auction 1 for every finite value of the precision parameter.

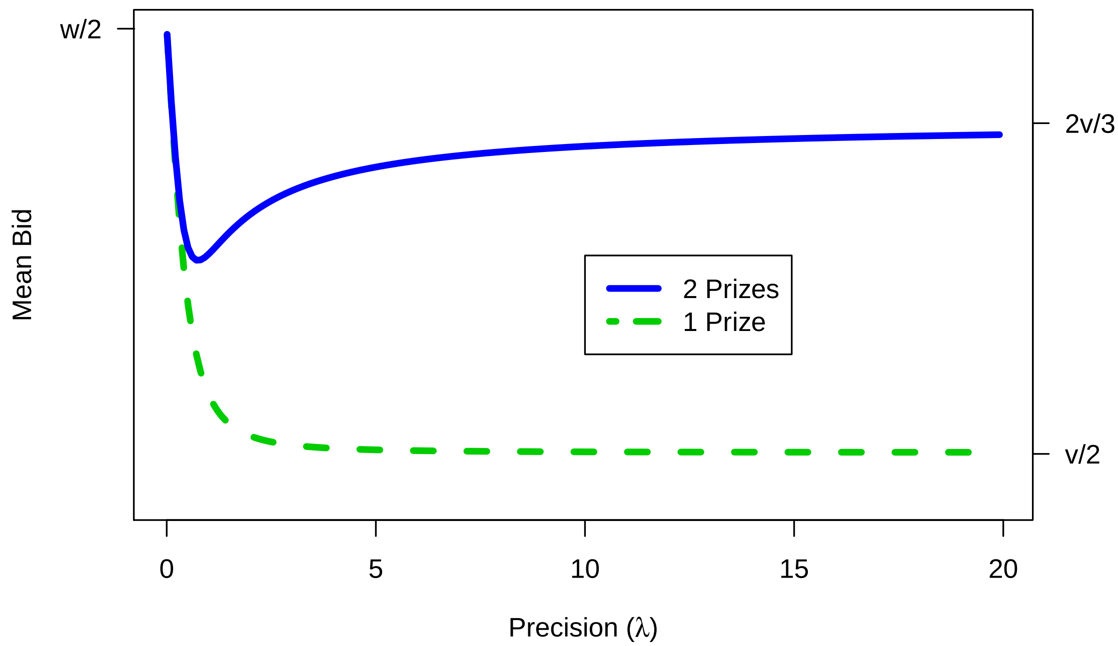

Figure 3: Mean bid under logit quantal response equilibrium.

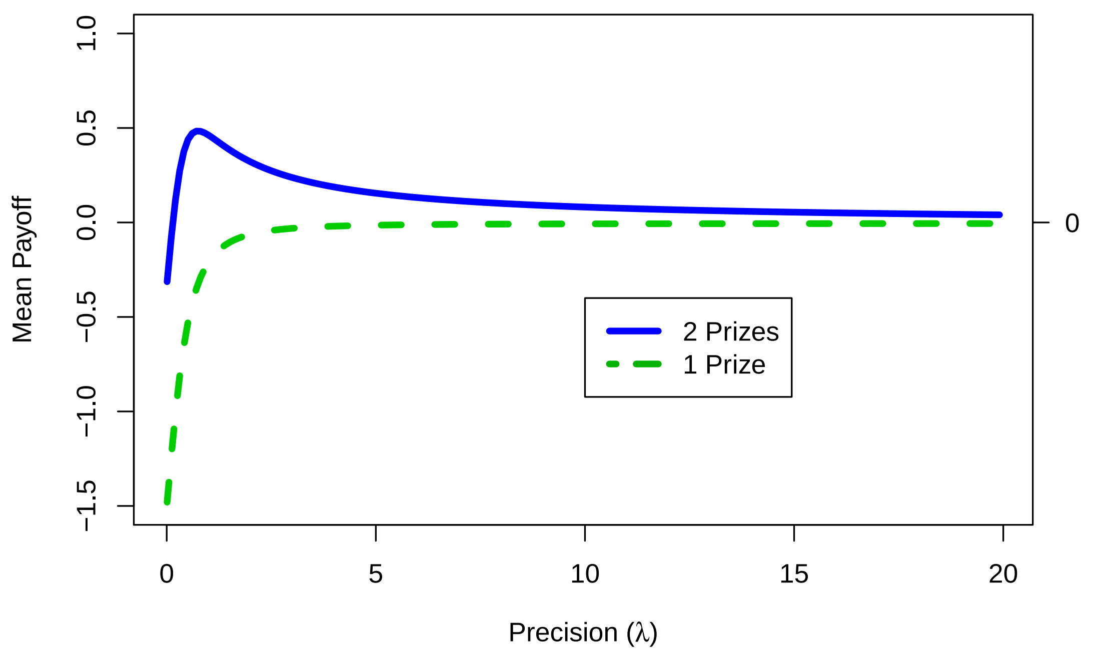

Figure 4: Mean payoff under logit quantal response equilibrium.

3.2 Evolutionary Models

Evolutionary models include both evolutionary stability criteria and evolutionary dynamics. Evolutionary stability criteria characterize the robustness of equilibria to small perturbations. Dynamic evolutionary models describe an adaptive process of behavioral adjustment in large populations of boundedly rational agents. When the stationary-points of these evolutionary dynamics correspond to equilibria, they often exhibit stability properties that correspond to the characterizations provided by evolutionary stability criteria.

3.2.1 Evolutionary Stability Criteria

To evaluate the stability of these Nash equilibria, we consider the 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

$$Q\left(Z\right)=\int D_F\pi(b|\Phi)Z(b)\,dZ(b) \tag{13}$$

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

$$Q\left(Z\right)=\int_0^w \left.\frac{\partial\pi\left(b|F\right)}{\partial F\left(b\right)}\right|_{F=\Phi} Z\left(b\right)dZ\left(b\right) \tag{14}$$

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

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

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

$$\int_0^w \tilde{\pi}\left(b|G\right) dZ(b) = \int_0^w \tilde{\pi}\left(b|G\right) dG(b) - \int_0^w \tilde{\pi}\left(b|G\right) d\Phi(b) = \tilde{\pi}(G|G) - \tilde{\pi}(\Phi|G). \tag{16}$$

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 approximate payoff to the alternate bidding strategy G is strictly greater than the approximate payoff to the equilibrium bidding strategy. Conversely, if $Q(Z)$ is strictly negative then the approximate payoff to the alternate bidding strategy G is strictly less than the approximate 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.

Proposition 2. *The Nash equilibrium for auction 2 is positive definite.*

Proof. See appendix on page 32. □

Proposition 2 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.

Proposition 3. *The Nash equilibrium for auction 1 is indefinite.*

Proof. See appendix on page 32. □

Proposition 3 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.

Proposition 4. *The Nash equilibrium for auction 1 is globally neutrally stable.*

Proof. See appendix on page 33. □

Proposition 4 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, proposition 2 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.

3.2.2 Evolutionary Dynamics

The logit dynamic describes the evolution of a mixed 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

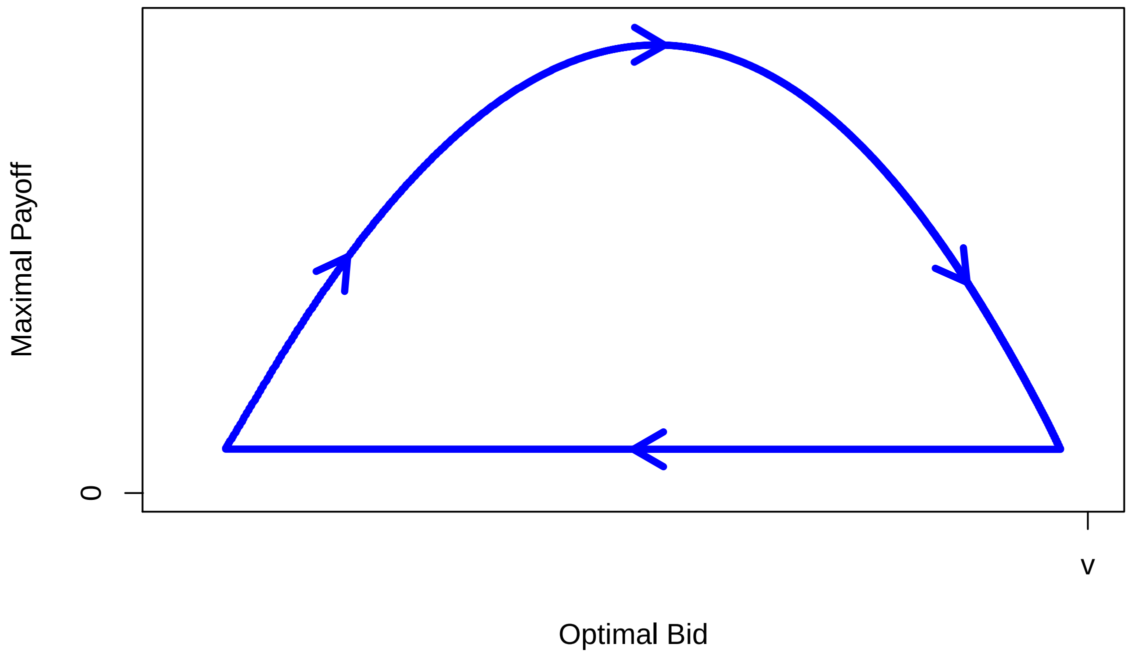

Figure 5: The limit cycle in auction 2 projected onto a two-dimensional space with the most profitable bid on the horizontal axis and the highest payoff on the vertical axis. The parameter values for this limit cycle are set to mirror the experimental setup for auction 2 and are given by $v = 7$, $w = 10$, and $\lambda = 10$.

that an adjusting agent will select the bid $b \in [0, w]$ is given by the logit response

$$\ell\left(b|F\right) = \frac{\exp \lambda\pi\left(b|F\right)}{\displaystyle\int_0^w \exp \lambda\pi\left(x|F\right) dx} \tag{18}$$

Hence the evolution of the bid distribution over time is governed by the differential equation

$$\dot{F}\left(b\right) = \mathcal{L}\left(b|F\right) - F\left(b\right) \text{ where } \mathcal{L}(b|F) = \int_0^b \ell(b|F) \tag{19}$$

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. Proposition 2 states that the unique mixed strategy Nash equilibrium of auction 2 is positive definite while proposition 4 states that the unique mixed strategy Nash equilibrium of auction 1 is globally neutrally stable. Hence 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 5 depicts the projection of this limit cycle onto the

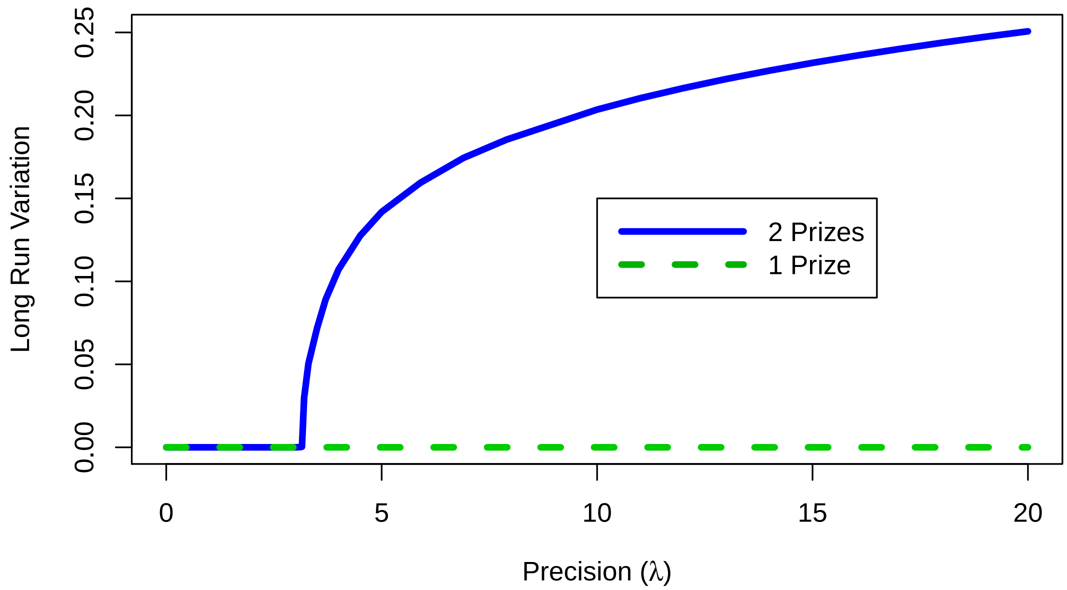

Figure 6: Long-run bid variation under the logit dynamic.

two dimensional space with optimal bids on the horizontal axis and optimal payoffs on the vertical axis.

Figure 6 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

$$\text{Var}(F) = \lim_{T\to\infty} \frac{1}{T} \int_0^T \left\| \bar{F} - F_t \right\| dt, \tag{20}$$

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 and the long-run variation is greater than zero if the distribution of bids converges to a limit cycle.

4 Experimental Design

The experiment utilized a between-subject design. Each session implemented one of the two all-pay auctions described in section 3. In treatment 1, two bidders compete over one-prize. In treatment 2, three bidders compete over two prizes. In both treatments, subjects were endowed with $w = \$10$ and competed for prizes with known value $v = \$7$. Each session had four 5 minute periods, during which subjects could adjust their bids as frequently as desired. When a subject clicked, her bid instantaneously changed to the level corresponding to the horizontal position of her mouse.[9] Subjects could use this interface to instantly select any bid in dollars and cents from \$0 to \$10.

To implement evolutionary 'playing the field' (Smith, 1982; Sandholm, 2010), we employ mean matching (e.g. Cason et al., 2014; Oprea et al., 2011) where each subject's instantaneous payoff is given by the expected value of her payoff from being randomly matched against the other subjects in her session. By the law of large numbers, mean-matching provides a superior approximation to truly continuous random matching than high frequency random matching. Consistent with the theoretical framework described in section 3, each subject effectively competes against the entire population of possible opponents. Subjects received continuous feedback regarding their mean-matching payoffs throughout each period. Bids and payoffs were recorded at a rate of ten times per second.[10] At the end of each session, subjects received the time average of their mean-matching payoff in addition to a fifteen dollar show-up payment.

The experiment also had two informational treatments, also implemented between subject. 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 could directly observe the current payoff to every possible bid. This paper focuses on the the auction treatment, while the informational treatments are addressed

[9]This instantaneous change is known as a "jump-adjustment:" the alternative "continuous-adjustment" has strategies gradually change upon subject input. Cason et al. (2014) employs each method in a separate treatment. Stephenson (2019) employs the latter method.

[10]Because payoffs are calculated ten times per second, one could interpret this as a finitely repeated game. This approximation of continuous time is common in the literature (see Cason et al., 2014; Oprea et al., 2011; Stephenson, 2019). Further, it is unlikely subjects had the cognitive ability or physical reflexes to make adjustments ten times per second, making the game effectively continuous from the subjects' standpoint.

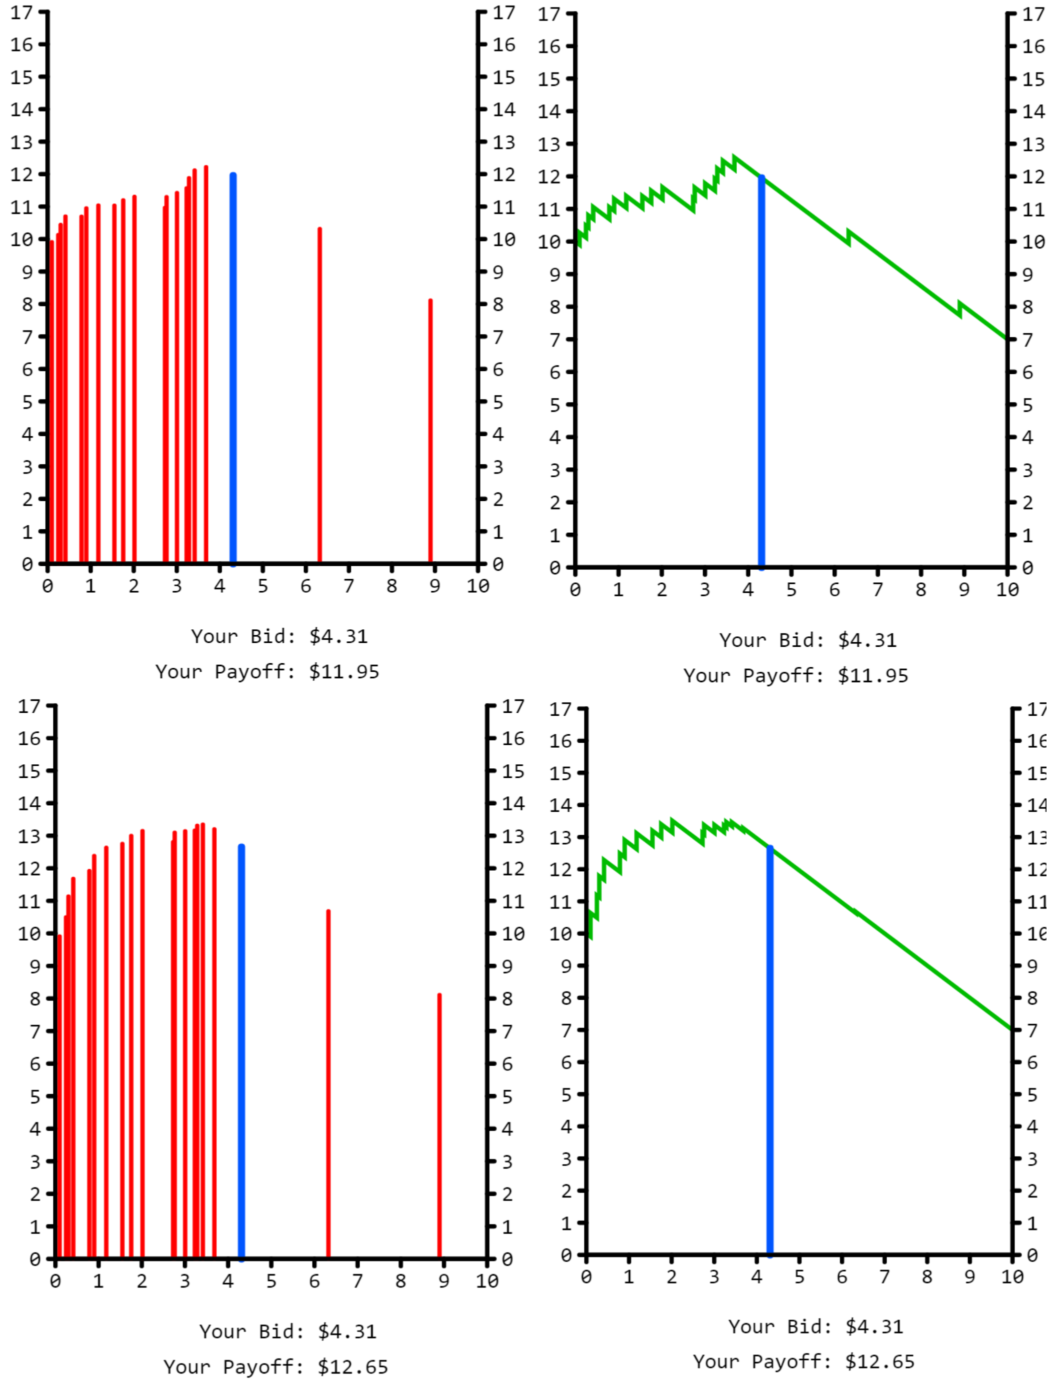

Figure 7: (a, top left) User interface under auction 1 with social information. (b, top right) user Interface under auction 1 with payoff information. (c, bottom left) User interface under auction 2 with social information. (d, bottom right) User interface under auction 2 with payoff information. The thick blue line denotes a subject's current bid and payoff. Under social information, the red lines represent the bids and payoffs of the other subjects. Under payoff information, the green line denotes the instantaneous payoffs associated with any bid.

	hypothesis 1: greater bids in auction 2 than 1	hypothesis 2: greater earnings in auction 2 than 1	hypothesis 3: greater instability in auction 2 than 1	hypothesis 4: clockwise cycles in auction 2
Nash equilibrium	yes	no	n/a	n/a
Logit QRE	yes	yes	n/a	n/a
Stability Criteria	n/a	n/a	yes	n/a
Logit Dynamic	yes	if $\lambda < 7^a$	yes	yes

[a] In unstable settings, the logit dynamics orbit the logit QRE in distribution space. This orbit is generally asymmetrical, so the time-average does not generally coincide with the logit QRE. If $\lambda < 7$, the logit dynamics predict higher earnings in auction 2 than auction 1. As shown in table A.2, maximum likelihood estimation finds $\lambda < 7$ for every subject under logit dynamics.

Table 1: Theoretical predictions of Nash equilibrium, logit quantal response equilibrium, stability criteria and logit dynamics. Nash and quantal response equilibrium models are designed to predict the fixed-point where strategic play converges and offer no predictions on the likelihood of convergence or the dynamics about a fixed point. Stability criteria predict the likelihood of convergence but make no prediction regarding comparative statics or disequilibrium dynamics.

in a separate paper (Stephenson and Brown, 2020).[11]

Figure 7 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, bids and payoffs of others are represented by red lines. In the payoff information treatment, the subject's instantaneous payoff function is represented by a green line.

Eight experimental sessions were conducted, four for each auction treatment condition. Each session was run with twenty subjects. All 160 subjects were recruited from the Texas A&M undergraduate population using an ORSEE database (Greiner, 2015). At the end of every session, each subject received the time average of their instantaneous payoff plus a five dollar show-up payment. Subject in auction 1 earned an average of \$15.47. Subjects in auction 2 earned an average of \$15.88. In equilibrium, average subject earnings were equal to \$15.00, so subjects received slightly above equilibrium earnings under both treatments. Every session lasted less than an hour, including instructions and payment procedures.

[11] None of the five outcome variables (see Table 2) differ substantially between information treatments, holding the auction constant (see Appendix Table A.1). Regression analysis (not provided) confirms the main results of this paper even in specifications that use information treatment as an explanatory variable.

5 Hypotheses

Section 3 provided the theoretical predictions of fixed-point models (i.e., Nash equilibrium, logit quantal response equilibrium), stability criteria, and evolutionary dynamics. While all of these models are applicable to the environment, they each motivate different types of hypotheses. Fixed-point models identify particular strategies which are often interpreted as describing where play will eventually converge. Such models are not designed to provide predictions regarding the likelihood of this convergence. In contrast, stability criteria focus on this latter issue with less focus on other details about the strategy profile. Dynamic evolutionary models explicitly describe a process where strategies change over time which may or may not lead to convergence in the long run. Table 1 provides a summary of the four theoretical models and their relation to the four hypotheses discussed in this section.

Both the Nash equilibrium and logit quantal response equilibrium predict a higher mean bid for auction 2 than auction 1. As illustrated by figure 1, the equilibrium bid distribution of auction 2 first-order stochastically dominates the equilibrium bid distribution of auction 1. Hence Nash equilibrium predicts that bidders will bid more aggressively in auction 2. As illustrated by figure 3, 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 every positive precision level.

Hypothesis 1. *Average bids will be greater 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 4, 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 under every finite precision level.

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

Proposition 4 indicates that the Nash equilibrium of auction 1 is globally neutrally stable. It does at least as well against any alternative strategy than the alternative does against itself. In contrast, proposition 2 indicates that the the Nash equilibrium strategy of auction 2 is positive definite. It does worse against any sufficiently close alternative strategy then the alternative does against itself. These theorems indicate that the Nash equilibrium of auction 2 is less stable than the Nash equilibrium of auction 1.

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

Figure 5 illustrates the clockwise limit cycles predicted by the logit dynamic for auction 2 in the two-dimensional space with the optimal bids on the horizontal axis and optimal payoffs on the vertical axis. Together with the local instability result from proposition 2, 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.*

6 Results

Table 2 provides summary statistics for all bid adjustments made by subjects over all sessions of the experiment. In total, subjects made 237,800 separate bid adjustments. Each individual subject made an average of 1486.25 bid adjustments over four, five-minute periods, roughly one bid adjustment every 0.8 seconds.[12] Bid adjustments were substantive, as each adjustment changed a bid by an average of \$0.89, roughly 9% of the strategy space. Subjects generally improved their payoff with each adjustment, as an adjustment increased a subject's instantaneous payoff by an average of \$0.41.

We now turn to testing the theoretical predictions of fixed-point models. Both Nash equilibrium and logit quantal response equilibrium predict a higher mean bid for auction 2 than auction 1.

Result 1. *Bids and payoffs were both higher in auction 2 than in auction 1. In both auctions, bids were lower than Nash predictions while payoffs were higher than Nash predictions.*

[12] Subjects in our experiment could instantaneously adjust their action. In Stephenson (2019) and some treatments of Cason et al. (2014), subjects made continuous adjustments to their action over time, so they lack comparable jumps that can be counted (see footnote 9). Of the studies with jump adjustments, Cason et al. (2014) and Oprea et al. (2011) do not mention the frequency of adjustments they observed. Stephenson (2020) observed an adjustment every 6.67 seconds in a 24-action, dominant strategy, school choice game.

	Overall	Auction 1	Auction 2
Average Bid	3.68 (1.79)	3.24 (1.79)	3.96 (1.76)
Dominated Bid (above 7)	0.02 (0.15)	0.01 (0.09)	0.03 (0.17)
Absolute Change in Bid	0.89 (1.08)	0.93 (1.07)	0.87 (1.09)
Payoff Change from Adjustment	0.41 (0.88)	0.30 (0.68)	0.48 (0.98)
Total Adjustments	237800	91105	146695
Total Minutes	160	80	80
Total Periods	32	16	16
Total Sessions	8	4	4
Total Subjects	160	80	80

Table 2: Summary statistics for bid adjustments. Standard deviations are provided in parentheses.

	Nash Equilibrium		Empirical 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)

Table 3: Nash predictions and corresponding period-level empirical outcomes by auction. Standard errors are provided in parentheses.

Table 3 summarizes five key empirical values at the period level alongside Nash equilibrium predictions for comparison. As predicted by fixed-point models, the average bid in auction 1 (3.038) was lower than the average bid in auction 2 (3.800). A permutation test, run at the session level (N=8), finds these differences to be significant at the 5% level (p-value ≈ 0.0286, two-tailed).[13] All sessions feature average bids lower than the corresponding Nash prediction.

Figures 8 (a) and (b) show histograms of subject mean bids in auctions 1 and 2, respectively. The distribution of mean subject bids loosely resemble a normal curve centered around a point slightly below the Nash prediction. Bids are typically higher in auction 2, consistent with our previous results. Most subjects bid below the mean Nash prediction on average. Only 28 and 8 subjects in auctions 1 and 2, respectively, have mean bid amounts higher than Nash levels. Figures 8 (c) and (d) show histograms for the overall distribution of subject mean bids in auctions 1 and 2, respectively. The overall distribution of bids exhibits bimodality in auction 2. However, in contrast to Ernst and Thöni (2013), the overall distribution of mean bids is roughly unimodal in auction 1.

Since earnings are given by winnings less bids, mean earnings are inversely related related to mean bids. Nash Equilibrium predicts mean earnings of 0 in both auction 1 and auction 2. In contrast, mean earnings exceeded 0 in all 8 auction sessions. Consistent with Hypothesis 2, average earnings were higher auction 2 (0.879) than in auction 1 (0.471) (p-value ≈ 0.0857, two-tailed). However, the predictions are not fully consistent with the logit quantal response equilibrium, which predicts positive earnings in auction 2 but *negative* earnings in auction 1.

Evolutionary models predict that auction 2 will exhibit greater instability than auction 1. 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-averaged distribution of bids over the entire period. The time-average of the Chebyshev distance between F_t and $\bar{F}$ serves as our measure of instability in the empirical bid distribution. We also compute F^*_A, the Nash equilibrium prediction for auction A and compute the time-average

[13]Unless otherwise noted, we provide comparisons of session level averages in our analysis. Using period-level averages instead would not qualitatively affect our results. If anything, using period-level treatment comparisons, whether parametric or non-parametric, would make the differences between treatments significant at lower thresholds.

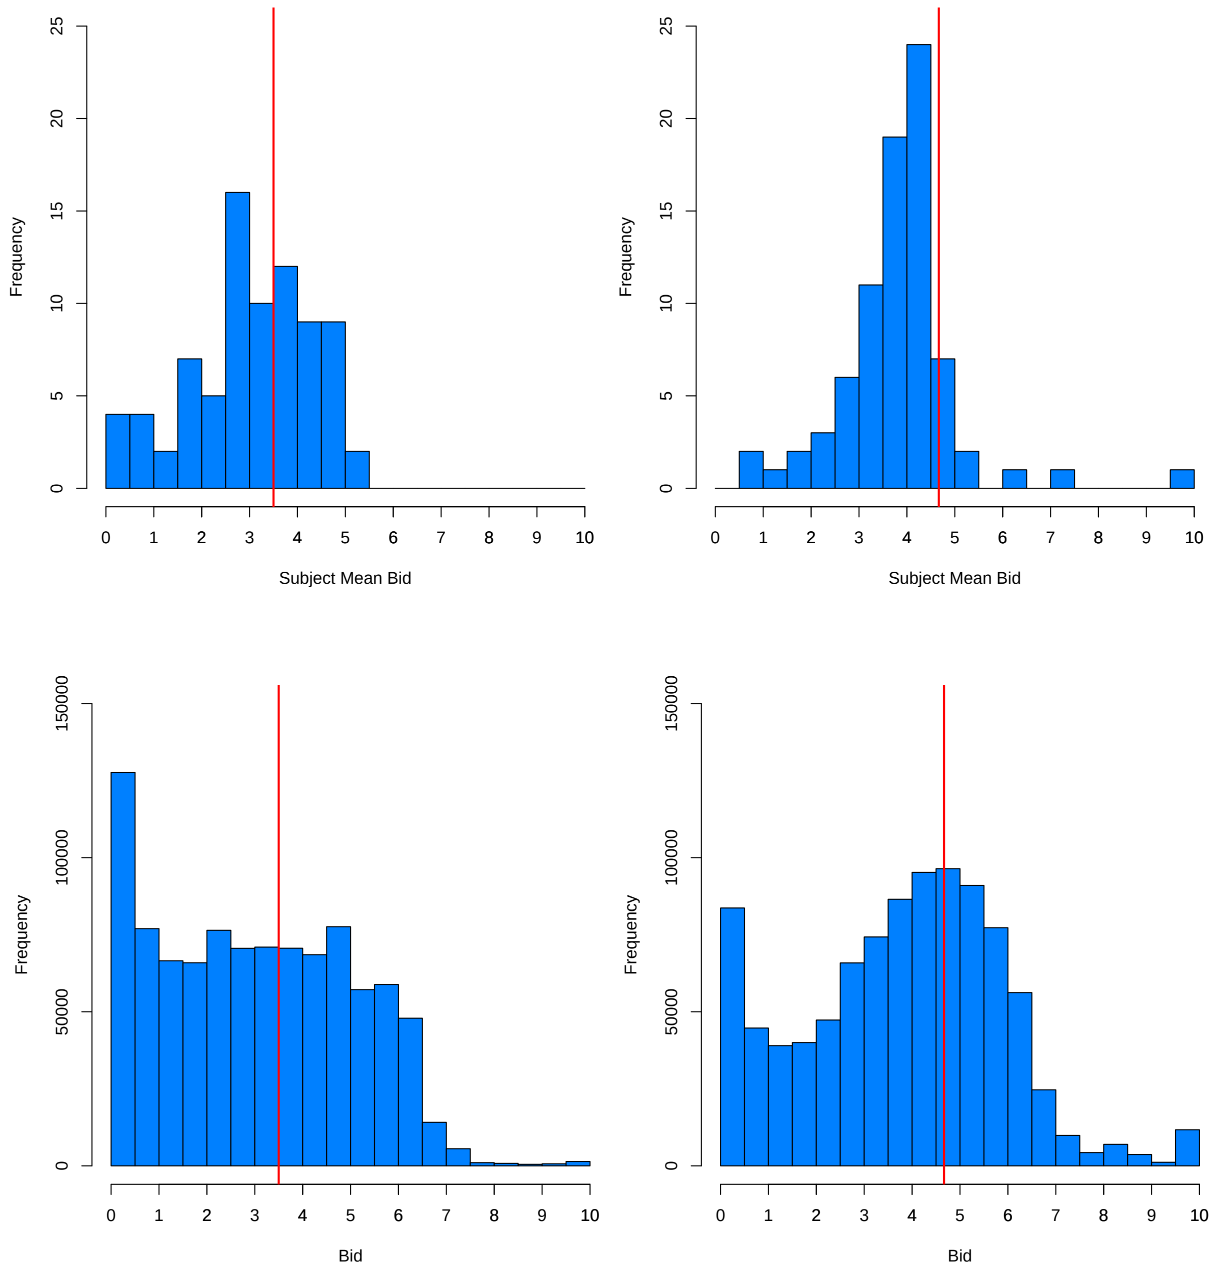

Figure 8: Histograms of observed bids. The first row illustrates the distribution of mean subject bids for auction 1 (a, left) and auction 2 (b, right). Mean subject bids were calculated for each subject over an entire 20 minute experimental session comprising 4 periods. The second row illustrates the overall distribution of bids for auction 1 (c, left) and auction 2 (d, right). The vertical red lines denote the mean Nash bid for each auction.

of the Chebyshev distance between F_t and F_A^* as a measure of deviation from equilibrium.

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

Table 3 provides the means and standard errors for this measure of instability over all 32 periods in the experiment, separated by treatment. Consistent with evolutionary stability criteria, sessions that implemented auction 2 exhibited greater instability (0.283) than sessions that implemented auction 1 (0.193) (p-value ≈ 0.0286, two-tailed). These results are consistent with the positive definiteness of auction 2 and the global neutral stability of auction 1. Both auctions exhibited substancial deviation from Nash equilibrium. This finding is consistent with Result 1 which showed that bids and payoffs were both different from their Nash equilibrium values. Deviation from equilibrium was greater in auction 2 (p-value ≈ 0.0286, two-tailed), which is consistent with the greater instability of auction 2.

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

Logit response dynamics predict clockwise cycling in both auctions. In contrast, neither Nash equilibrium nor quantal response equilibrium is equipped to address the presence of cyclical behavior. In auction 1, these cycles are predicted to converge on a logit quantal response equilibrium. In auction 2, these cycles are predicted to converge on a stable limit cycle orbiting the logit quantal response equilibrium. To measure these cycles, we project the infinite dimensional distribution space onto a more manageable two-dimensional space with optimal bids on the horizontal axis and optimal payoffs on the vertical axis. We can easily observe a great deal of persistent clockwise cycling in this two-dimensional space. Figure 9(a) provides one example of such a cycle that occurred during period 2 of session 4.

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

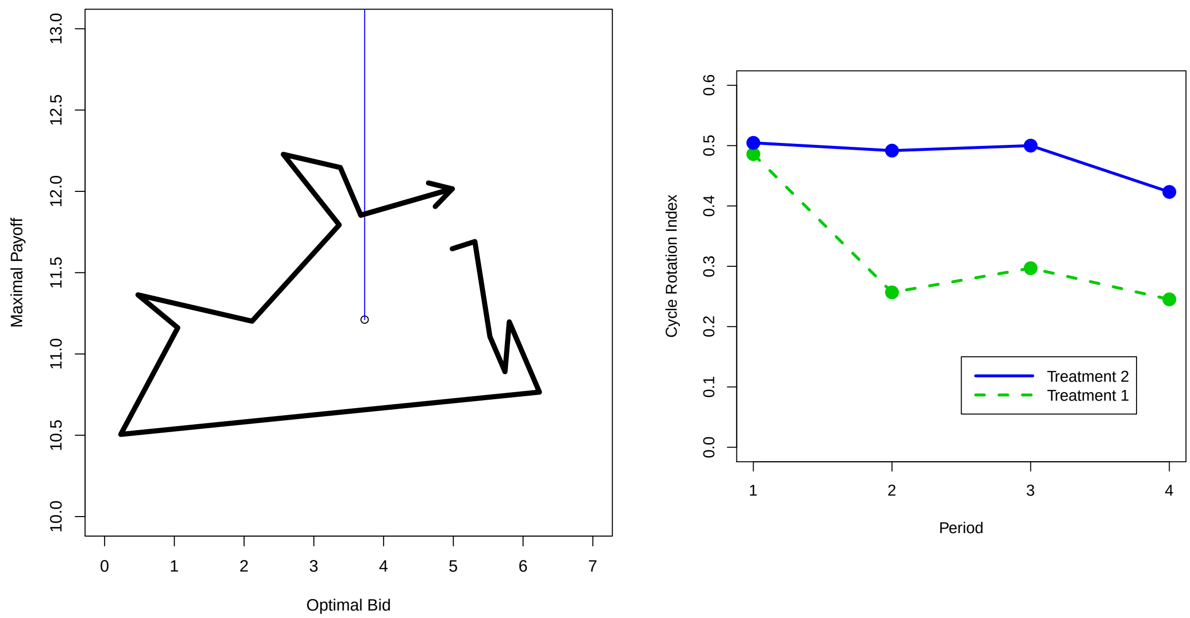

Figure 9: (a, left) An example cycle on the two-dimensional graph of optimal bid and maximal payoff. Observed from seconds 16 to 21 of period 2, session 4. (b, right) Mean cycle-rotation index by auction and period.

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 9(a). Next we calculate the number of clockwise rotations and the number of counterclockwise rotations. A rotation is said to occur when the bid 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. 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 we infer the distribution of bids exhibits significant cyclical patterns.

Logit dynamics predict clockwise cycling in both auction 1 and auction 2. In auction 1, these cycles converge on the logit QRE. However, in auction 2, these cycles exhibit persistent non-convergence. Accordingly, logit dynamics predict that long-run behavior will coincide with logit QRE in auction 1, but not in auction 2. Consistent with these predictions, maximum likelihood estimation[14] finds that 85% of subjects in treatment 2 are better explained by the logit dynamic than by the logit QRE, while only 52% of subjects in auction 1 are better explained by the logit dynamic than by the logit QRE.

Table 3 provides the mean and standard deviation of the Cycle-Rotation Index for auction 1 and auction 2 at the period level. All eight sessions feature average CRI's above 0, indicating pronounced clockwise cycling in both treatments. All eight values are also below 1, the theoretical prediction of perfect clockwise cycling, indicating the presence of behavioral noise. The cycle rotation index is higher in auction 2, consistent with the theoretically predicted convergence of auction 1 (p-value ≈ 0.0857, two-tailed). Figure 9(b) provides CRI averages by period in all eight sessions. The CRI measure is above 0 in all periods, indicating strong clockwise cycling. The CRI decreases in later periods of auction 1, consistent with the predicted long run convergence in auction 1. The CRI remains relatively constant in auction 2, consistent with the predicted persistence of limit cycles in auction 2.

7 Discussion

This paper connects two strands of literature: experimental investigations of the all-pay auction and experimental studies of evolutionary dynamics in continuous time. In both areas this work departs from convention. It is the first examination of the all-pay auction in continuous time. It also marks the first continuous-time experimental examination of mixed strategy equilibria over a continuous strategy space, rather than two or three pure strategies. In contrast to previous discrete-time investigations of the all-pay auction, we observe persistent underbidding relative to the Nash predictions. Consistent with evolutionary models, we observe persistent cyclical behavior in both treatments and greater stability in our first treatment.

Most previous experimental studies of the all-pay auction find overbidding relative to

[14] See table A.2 for a detailed description of the maximum likelihood estimates.

Nash equilibrium. All of these were conducted in discrete time with random matching or fixed groups. Mean-matching in continuous-time mirrors the structure of evolutionary models and provides subjects with far more experience than conventional discrete time protocols. Accordingly, our findings suggest that overbidding in all-pay auctions diminishes with experience. Indeed, previous literature that examined a large number of discrete periods noted a declining trend of overbidding with experience (Davis and Reilly, 1998; Lugovskyy et al., 2010). While inexperience may indeed contribute to overbidding in all-pay auctions, other explanations also have empirical support. Utility of winning and competitive preferences might explain overbidding in a way that would persist with experience. The use of mean-matching in continuous-time may increase the salience of payoffs while diminishing the salience of winning. These experimental protocols may have suppressed competitive preferences that subjects would have exhibited under conventional matching in discrete-time. Agents in the field who occasionally participate in an auction may experience utility of winning that subjects in our experiment did not. Conversely, agents in the field who routinely participate in many auctions may seek to maximize payoffs in a way that subjects in conventional discrete-time experiments do not. It remains for the reader (or future experimenter) to determine which experimental protocols are most useful for the questions they aim to investigate.

Previous experimental studies investigating mixed-strategy equilibria in continuous-time have considered games with a small number of pure strategies. Stephenson (2019) and Cason et al. (2014) study mixed-strategies over two or three actions, respectively. These experiments allowed subjects to directly select probability distributions over actions. In contrast, the all pay auction features a continuous strategy space which we discretize into 1001 actions. Graphically illustrating the 1000 dimensional space of mixed strategies was infeasible, so our experimental interface allowed subjects to directly select bids. As in evolutionary models, mixed strategies occur as distributions of actions over many subjects rather than randomization over actions by an individual subject. Our experiment involved only a finite number of subjects, so it only approximates the large population limit described by evolutionary models. The slight underbidding we observe may result from these limitations.[15]

Comparing the observed behavior across treatments, we find that the distribution of bids exhibited greater stability over time in auction 1 than in auction 2. We also observe persistent cyclical dynamics in both auctions. These observations cannot be explained by fixed-point models, but they are consistent with dynamic evolutionary models. These results suggest that evolutionary models can inform practitioners in a wide variety of concrete settings outside the laboratory about whether long-run aggregate behavior is likely to approach equilibrium predictions.

Consider the designer of a contest structure to allocate grant funding. Depending on their policy goals, the designer may prefer some outcomes over others. To evaluate potential contest structures, the designer might employ models that identify which outcomes are likely to occur under each contest structure. A designer who only considers conventional fixed-point models may incorrectly assess the desirability of a given contest structure because they lack important information provided by evolutionary models. Several distinct contest structures may yield identical equilibrium predictions, but differ in evolutionary predictions. A policymaker who fails to consider evolutionary models may select a contest structure that has desirable equilibrium properties but induces undesirable non-convergent behavior. Further, it is conceivable that policymakers who prefer more stable and predictable outcomes might prefer contest structures that induce reliable convergence to a particular equilibrium. It is straightforward to apply similar arguments to other concrete field settings (refer to the introduction for many examples).

While additional research is still required to establish the generality of these results in a wider class of strategic environments, initial work is promising. A variety of papers now present evidence for the usefulness of adaptive models in predicting convergence to equilibrium (i.e., Benndorf et al., 2016; Bigoni et al., 2015; Cason et al., 2014, 2020; Oprea et al., 2011; Stephenson, 2019). We encourage more experimental, theoretical, and applied research in this burgeoning area.

[15]If it had been feasible to clearly depict a payoff distribution over the 1000 dimensional space of mixed strategies, then we might have asked subjects to directly select mixed strategies. Alternatively, if each session involved thousands of subjects, then we might have obtained a closer approximation of the large population limit. In such experiments, time-averaged bids might have been even closer to the Nash predictions.

References

Abreu, D. and Sethi, R. (2003). Evolutionary stability in a reputational model of bargaining. *Games and Economic Behavior*, 44(2):195–216.

Alós-Ferrer, C., Ania, A. B., and Schenk-Hoppé, K. R. (2000). An evolutionary model of bertrand oligopoly. *Games and Economic Behavior*, 33(1):1–19.

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.

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.

Benndorf, V., Martinez-Martinez, I., and Normann, H.-T. (2016). Equilibrium selection with coupled populations in hawk–dove games: Theory and experiment in continuous time. *Journal of Economic Theory*, 165:472–486.

Bigoni, M., Casari, M., Skrzypacz, A., and Spagnolo, G. (2015). Time Horizon and Cooperation in Continuous Time. *Econometrica*, 83:587–616.

Brown, A. and Stephenson, D. (2019). *Handbook of Experimental Game Theory*, chapter Games with Continuous-Time Experimental Protocols. Forthcoming.

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.

Cason, T. N., Friedman, D., and Hopkins, E. (2020). An experimental investigation of price dispersion and cycles.

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

Davis, D. D. and Holt, C. A. (1993). *Experimental economics*. Princeton university press.

Davis, D. D. and Reilly, R. J. (1998). Do too many cooks always spoil the stew? an experimental analysis of rent-seeking and the role of a strategic buyer. *Public Choice*, 95(1-2):89–115.

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.

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.

Hens, T. and Schenk-Hoppé, K. R. (2005). Evolutionary stability of portfolio rules in incomplete markets. *Journal of mathematical economics*, 41(1-2):43–66.

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.

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.

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.

Potters, J., De Vries, C. G., and Van Winden, F. (1998). An experimental examination of rational rent-seeking. *European Journal of Political Economy*, 14(4):783–800.

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

Smith, J. M. (1982). *Evolution and the Theory of Games*. Cambridge university press.

Stephenson, D. (2019). Coordination and evolutionary dynamics: When are evolutionary models reliable? *Games and Economic Behavior*, 113:381–395.

Stephenson, D. (2020). Assignment feedback in school choice mechanisms. *Working Paper*.

Stephenson, D. G. and Brown, A. L. (2020). Characterizing persistent disequilibrium dynamics: Imitation or optimization?

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

A Proofs

Proof of Proposition 2. 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

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

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 (13) is given by

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

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

Proof of Proposition 3. 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

$$\left.\frac{\partial\pi\left(b|F\right)}{\partial F\left(b\right)}\right|_{F=\Phi} = v \tag{A.3}$$

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 (14) can be written as

$$Q(Z) = v\int_0^w Z\left(b\right) dZ\left(b\right) = \frac{1}{2}vZ(w)^2 \tag{A.4}$$

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

Proof of Proposition 4. 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

$$\begin{aligned}
\int b\, dF(b) &= \int_{b=0}^{b=w} \int_{x=0}^{x=b} dx\, dF(b) \\
&= \int_{x=0}^{x=w} \int_{b=x}^{b=w} dF(b)\, dx \\
&= \int_0^w [1 - F(x)]\, dx \qquad \text{(A.5)}
\end{aligned}$$

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

$$\begin{aligned}
\pi(\Phi|F) &= v \int F(b)\, d\Phi(b) - \int b\, d\Phi(b) \\
&= \int_0^v F(b)\, db - \frac{v}{2} \qquad \text{(A.6)}
\end{aligned}$$

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

$$\begin{aligned}
\pi(F|F) &= v \int F(b)\, dF(b) - \int b\, dF(b) \\
&= \frac{v}{2} - \int b\, dF(b) \qquad \text{since } F(X) \sim U[0,1] \text{ for } X \sim F(x) \\
&= \frac{v}{2} - \int_0^w [1 - F(b)]\, db \qquad \text{since } \int b\, dF(b) = \int_0^w [1 - F(b)]\, db \\
&\leq \frac{v}{2} - \int_0^v [1 - F(b)]\, db \\
&= \frac{v}{2} - v + \int_0^v F(b)\, db \\
&= \int_0^v F(b)\, db - \frac{v}{2} \\
&= \pi(\Phi|F) \qquad \text{(A.7)}
\end{aligned}$$

Thus the equilibrium strategy Φ does weakly better against the non-equilibrium strategy F then the non-equilibrium strategy F does against itself. □

Proof of Proposition 1. 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

$$\begin{aligned}
\frac{dL}{db} &= \frac{\exp\left(\lambda\left(2vL(b) - vL(b)^2 - b\right)\right)}{\int_0^w \exp\left(\lambda\left(2vL(y) - vL(y)^2 - y\right)\right) dy} \\
C\, dL &= \exp\left(\lambda\left(2vL - vL^2 - b\right)\right) db \\
C \exp\left(\lambda v\left(L^2 - 2L\right)\right) dL &= \exp\left(-\lambda b\right) db \\
C \exp(-\lambda v) \int \exp\left(\lambda v \left(1 - L\right)^2\right) dL &= \int \exp\left(-\lambda b\right) db \\
\operatorname{erfi}\left(\sqrt{\lambda v}\left(1 - L\right)\right) &= C_1 - C_2 \exp\left(-\lambda b\right) \\
L\left(b\right) &= 1 - \frac{1}{\sqrt{\lambda v}} \operatorname{erfi}^{-1}\left(C_1 - C_2 \exp\left(-\lambda b\right)\right) \qquad \text{(A.8)}
\end{aligned}$$

Where erfi denotes the imaginary error function and is given by $\operatorname{erfi}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(x^2) dx$. 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

$$\begin{aligned}
C_1 &= \frac{\operatorname{erfi}\left(\sqrt{\lambda v}\right)}{\exp\left(-\lambda w\right) - 1} + \operatorname{erfi}\left(\sqrt{\lambda v}\right) \\
C_2 &= \frac{\operatorname{erfi}\left(\sqrt{\lambda v}\right)}{\exp\left(-\lambda w\right) - 1} \qquad \text{(A.9)}
\end{aligned}$$

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

$$L\left(b\right) = 1 - \frac{1}{\sqrt{\lambda v}} \operatorname{erfi}^{-1}\left(\left[1 - \frac{1 - \exp\left(-\lambda b\right)}{1 - \exp\left(-\lambda w\right)}\right] \operatorname{erfi}\left(\sqrt{\lambda v}\right)\right) \qquad \text{(A.10)}$$

□

B Additional Tables

Table A.1: Session level averages for five main outcome variables. Standard errors, based on the four periods that comprise each session, are provided in parentheses

Session	Auction Treatment (no. of prizes)	info type	Mean Bid Amount)	Mean Payoffs	Deviation from Time-Averaged Mean	Deviation from Nash Equilibrium	Cycle-Rotation Index)
1	2	payoff	3.772 (0.162)	0.908 (0.164)	0.283 (0.022)	0.391 (0.044)	0.561 (0.043)
2	1	payoff	2.985 (0.113)	0.526 (0.113)	0.216 (0.010)	0.253 (0.012)	0.450 (0.079)
3	2	social	4.013 (0.259)	0.663 (0.259)	0.275 (0.011)	0.352 (0.028)	0.418 (0.033)
4	1	social	3.047 (0.057)	0.464 (0.058)	0.167 (0.013)	0.207 (0.005)	0.232 (0.054)
5	1	social	3.360 (0.087)	0.149 (0.086)	0.190 (0.008)	0.204 (0.005)	0.256 (0.057)
6	2	social	3.470 (0.167)	1.212 (0.166)	0.285 (0.021)	0.446 (0.019)	0.412 (0.055)
7	1	payoff	2.762 (0.063)	0.746 (0.062)	0.198 (0.011)	0.264 (0.003)	0.347 (0.062)
8	2	payoff	3.944 (0.115)	0.734 (0.114)	0.289 (0.022)	0.365 (0.008)	0.529 (0.061)

Table A.2: Maximum likelihood estimates. The precision parameter λ was estimated separately for each subject and each model.

				Log Likelihood		Estimated λ	
Session	Treatment	Subject	Adjustments	QRE	Logit Dynamic	QRE	Logit Dynamic
1	2	1	343	-14534.6	-13796.3	1.963215	0.996535
1	2	2	241	-10981.8	-9627.35	90.50562	0.968458
1	2	3	1866	-83741.1	-74238.9	75.15213	1.175507
1	2	4	2651	-120984	-106401	63.94754	1.161345
1	2	5	1424	-62930	-53493	75.38627	1.51164
1	2	6	4610	-210352	-188595	68.06483	1.025568
1	2	7	1015	-43521.6	-41737.1	0.866019	0.908018
1	2	8	1047	-45224.2	-44030.3	0.783802	0.861141
1	2	9	91	-3980.22	-3971.58	0.83077	0.561326
1	2	10	865	-37706	-36886.7	0.867565	0.758112

1	2	11	4206	-178294	-159793	0.919223	1.484914
1	2	12	3793	-166876	-158997	0.633957	0.87064
1	2	13	2289	-103642	-89013.7	89.49573	1.340342
1	2	14	977	-42418.5	-40802.1	0.784853	0.82635
1	2	15	1860	-82746.8	-68081.1	99.11586	1.565908
1	2	16	3819	-168709	-143541	76.0989	1.5124
1	2	17	693	-29749.9	-29541.5	1.057556	0.692377
1	2	18	1882	-82221.5	-81134.8	0.932834	0.668187
1	2	19	2207	-99366.1	-93760.6	65.48425	0.781501
1	2	20	1442	-63583.4	-57514.8	99.96539	1.218844
2	1	1	815	-34704.5	-36103.5	4.199217	0.644632
2	1	2	1166	-49639.2	-48568.4	2.901665	1.154651
2	1	3	835	-35534.4	-34033.1	91.45752	1.266575
2	1	4	1312	-56190.4	-54118.7	2.776578	1.226842
2	1	5	1021	-43449.8	-46229.3	38.19639	0.404926
2	1	6	643	-27483	-27252.3	11.28833	0.944863
2	1	7	257	-10936.9	-10960	75.51809	0.846774
2	1	8	2575	-110195	-112468	4.5569	0.7479
2	1	9	829	-35279	-37055.3	38.1966	0.543453
2	1	10	3246	-138162	-135621	12.27084	1.137972
2	1	11	686	-29193.5	-30130.7	76.3932	0.709913
2	1	12	686	-29923.6	-29287.2	2.590952	0.987081
2	1	13	1214	-51677	-48095.6	100	1.498592
2	1	14	1134	-50014.1	-48757.2	1.854483	0.859379
2	1	15	1428	-60835.9	-57818.5	7.123663	1.364025
2	1	16	1793	-76509.6	-72290.1	3.835422	1.426551
2	1	17	2861	-121968	-115646	3.604319	1.408393
2	1	18	945	-40375.8	-38227.2	3.631381	1.339872
2	1	19	3015	-128478	-125796	5.779932	1.182146
2	1	20	824	-35099.2	-33858.9	3.560629	1.222323
3	2	1	108	-4971.59	-4976.25	0.099239	0.111046
3	2	2	859	-36657.8	-37906.2	1.057556	0.570377
3	2	3	2178	-96569.6	-91157.2	77.52298	0.938038
3	2	4	1713	-75374.9	-73400.2	1.208891	0.758241
3	2	5	322	-14766.7	-14839.2	0.331603	0.108061
3	2	6	2407	-106802	-105884	0.781625	0.565653

3	2	7	2008	-84655.2	-81952	1.235869	1.119868
3	2	8	3323	-138609	-138804	2.273637	0.927364
3	2	9	667	-28296.5	-29364.5	1.361384	0.658704
3	2	10	4288	-183606	-168790	1.001847	1.372997
3	2	11	1775	-81918.4	-81917.4	3.44E-08	0.007485
3	2	12	1775	-81947.2	-69883.5	69.69383	1.280999
3	2	13	1207	-56367.5	-46324.6	60.17603	1.322825
3	2	14	1177	-50304.9	-46704.1	1.111115	1.296121
3	2	15	788	-33748.6	-30166	0.917922	1.457755
3	2	16	3176	-133274	-134770	1.382532	0.834053
3	2	17	695	-29767.7	-28455.5	1.154172	1.030427
3	2	18	2109	-90432.2	-85060.9	0.979474	1.150862
3	2	19	3704	-152250	-143697	1.661221	1.424304
3	2	20	2894	-123534	-119392	0.967682	1.022117
4	1	1	526	-22384.5	-24034.6	39.79382	0.313904
4	1	2	2181	-92938.3	-86848.3	6.115076	1.970737
4	1	3	1899	-83804.1	-83042.4	1.859269	0.838853
4	1	4	3819	-162739	-159815	5.568746	1.396205
4	1	5	2619	-112196	-112197	4.484998	1.095484
4	1	6	1753	-74981.4	-77285.4	3.984234	0.784049
4	1	7	1448	-61650.7	-63211.9	5.191696	0.942568
4	1	8	114	-4851.4	-5184.26	38.1966	0.386706
4	1	9	258	-11023.4	-11519	7.124152	0.628839
4	1	10	1422	-60665.2	-58185	6.183888	1.51685
4	1	11	647	-28464.8	-29493	1.976165	0.337191
4	1	12	134	-5894.35	-6147.46	0.91364	0.228161
4	1	13	402	-18451	-18450.9	0.668469	0.210904
4	1	14	196	-8401.4	-8901.38	0.985317	0.405714
4	1	15	815	-34719.2	-36654.9	2.771622	0.562126
4	1	16	49	-2123.87	-2147.7	2.407135	0.756388
4	1	17	2389	-101784	-102248	6.879742	1.155182
4	1	18	1059	-45090.1	-44218.8	8.053346	1.411199
4	1	19	396	-16956.7	-15751.8	5.675449	1.703004
4	1	20	284	-12180.9	-12697.3	2.326742	0.58869
5	1	1	1207	-51908.9	-51971.4	3.454045	1.110179
5	1	2	3059	-132025	-129068	3.968423	1.342237

5	1	3	1071	-45742.4	-44498.1	5.569852	1.459194
5	1	4	1415	-60602.8	-59157.5	3.265753	1.443999
5	1	5	1173	-50117.9	-52147.3	10.79762	0.711548
5	1	6	678	-28980	-29028.2	3.616434	1.156091
5	1	7	443	-19002.2	-18179.6	3.623336	1.617607
5	1	8	1505	-64762.6	-65450.5	2.641337	1.002699
5	1	9	713	-30835	-31880	6.79236	0.67447
5	1	10	4666	-198887	-208196	2.486891	0.690922
5	1	11	847	-36249.3	-32917.6	2.398916	2.160517
5	1	12	385	-16589.1	-17187.5	2.116019	0.677124
5	1	13	1125	-50101.2	-50158.2	1.454903	0.670774
5	1	14	273	-11924.7	-12203	1.400592	0.681927
5	1	15	21	-935.283	-967.998	0.791703	0.10198
5	1	16	37	-1581.51	-1707.44	99.99999	0.029255
5	1	17	467	-20257.2	-19776.9	3.008538	1.139379
5	1	18	732	-31684.9	-31361.1	5.335665	1.122488
5	1	19	1120	-47952.9	-49505.2	4.003602	0.81686
5	1	20	614	-26901.5	-28067.2	1.233643	0.3481
6	2	1	735	-31539.4	-29424.4	1.105023	1.117953
6	2	2	1493	-59412.2	-62037.2	1.663968	0.902007
6	2	3	1812	-76287.6	-69608	1.005909	1.451146
6	2	4	996	-42407.1	-38258.5	0.880398	1.432364
6	2	5	531	-23633.6	-23271.5	0.751594	0.564811
6	2	6	2438	-105468	-95862.7	0.810325	1.251926
6	2	7	1210	-50211.9	-44731.7	1.288546	1.600027
6	2	8	317	-13711.9	-13064.2	0.889258	0.951395
6	2	9	1297	-56429.2	-55628.7	0.818762	0.714589
6	2	10	436	-18535.8	-17016	0.814714	1.197823
6	2	11	3141	-138166	-136371	0.575758	0.627416
6	2	12	272	-12354.3	-11565.7	99.53492	0.716011
6	2	13	3745	-154166	-146216	1.182526	1.314665
6	2	14	1254	-54264.3	-50667	0.757358	1.090103
6	2	15	1304	-59126.5	-47063.2	77.54775	1.802036
6	2	16	2223	-93761.8	-89115.4	0.979683	1.141275
6	2	17	418	-17993.7	-16226.5	1.073568	1.194937
6	2	18	287	-12967.4	-12914	0.353848	0.369811

6	2	19	1469	-63537.7	-61475.9	0.884735	0.852802
6	2	20	1221	-51097.9	-48322.7	0.967682	1.224118
7	1	1	86	-3659.83	-3527.94	76.33327	1.150586
7	1	2	1453	-61841	-62753	91.15481	0.870367
7	1	3	635	-27023.1	-24074.9	76.3932	1.993097
7	1	4	938	-40079.4	-39796.1	3.69854	0.979546
7	1	5	664	-28658.3	-28850.7	2.38338	0.784076
7	1	6	1375	-58541.5	-53865	6.604801	1.740598
7	1	7	778	-33108.7	-31447.4	99.99989	1.447663
7	1	8	376	-16001.1	-16239.2	76.39337	0.873182
7	1	9	887	-37803	-37221.2	2.793055	1.116807
7	1	10	21	-893.679	-959.595	38.18243	0.293264
7	1	11	1349	-57651.4	-54963.9	3.386625	1.367242
7	1	12	292	-12426.4	-13249.4	38.19263	0.366788
7	1	13	1163	-49680.4	-48988.4	1.682929	1.076938
7	1	14	1963	-83614.6	-77268.9	4.833681	1.735934
7	1	15	818	-34922.4	-33781.7	5.861298	1.249318
7	1	16	1514	-64863.7	-65332	2.671415	0.852364
7	1	17	2491	-106261	-102626	3.22747	1.315209
7	1	18	867	-36923.7	-35433.4	7.311765	1.357056
7	1	19	1191	-51289	-48688.2	3.191925	1.372127
7	1	20	998	-42650.8	-39555.8	3.46028	1.601074
8	2	1	3803	-158301	-155782	1.164843	0.972915
8	2	2	4246	-177405	-165843	1.265171	1.23615
8	2	3	3264	-148626	-136295	77.4455	0.857876
8	2	4	2239	-94005.8	-98041.2	1.021509	0.560308
8	2	5	4663	-199067	-201451	0.911364	0.653619
8	2	6	1605	-74072.7	-74072.7	9.3E-09	9.3E-09
8	2	7	1319	-53710.4	-54441.7	2.281018	0.902918
8	2	8	1158	-50826.3	-48861	0.760287	0.76323
8	2	9	2944	-131288	-112931	65.78885	1.376602
8	2	10	577	-26127.7	-25003.1	0.307964	0.612133
8	2	11	1584	-68839.6	-59716.7	0.784149	1.473457
8	2	12	2490	-103892	-108001	1.371473	0.628896
8	2	13	1399	-63504.4	-54213.1	99.42059	1.294673
8	2	14	3914	-167969	-163116	0.805177	0.858727

8	2	15	1557	-68838.1	-67635.6	0.61816	0.591616
8	2	16	3520	-149584	-145596	1.294117	0.902779
8	2	17	832	-35507.7	-34378.6	1.270234	0.880409
8	2	18	1206	-51301.5	-48406.2	1.322736	1.064098
8	2	19	1755	-80595.3	-68707.1	63.92637	1.169551
8	2	20	1527	-68757.2	-61666.2	60.29398	1.006837