Convergence Rates in Resource Allocation Games

Daniel Stephenson*
stephensod@vcu.edu

Abstract

This paper experimentally investigates settings where agents allocate limited resources to compete for divisible complementary factors. When the success function is highly responsive to resource allocations, competitors have stronger incentives to best respond, but adaptive models predict slower convergence to equilibrium due to nonequilibrium incentives. In contrast, the unique Nash equilibrium allocations are shown to be proportional to prize values and do not depend on the responsiveness of the success function. To test these predictions, the experimental design varies both the prize values and the responsiveness of the success function independently across treatment conditions. Consistent with adaptive predictions, less responsive success functions produced faster convergence to equilibrium, suggesting that nonequilibrium incentives can influence the rate of convergence.

Keywords: Resource Allocation, Convergence, Responsiveness, Evolution

JEL Classification: C73, C92, D83, D91

*Virginia Commonwealth University, 907 Floyd Ave, Richmond, VA 23284

1 Introduction

To predict the impact of hypothetical policy interventions, economists often consider how they would affect equilibrium predictions. Yet the transition from a pre-intervention equilibrium to a post-intervention equilibrium may be gradual if economic agents are boundedly rational. This notion of gradual convergence to Nash equilibrium goes back as least as far as Nash's (1950) mass action[1] interpretation of equilibrium describing agents who lack "the ability and inclination to go though any complex reasoning process," but instead "accumulate empirical information on the relative advantages of the various pure strategies at their disposal." Adaptive models formalize this notion by explicitly describing how boundedly rational agents adjust their behavior over time.

This paper experimentally tests adaptive predictions about the rate of convergence to equilibrium in settings where agents allocate limited resources to compete for divisible complementary factors. Understanding the rate of convergence is an important part of evaluating of potential policy interventions since we are all dead in the long run, as noted by Keynes (1923). In empirical settings, strategic interaction frequently involves competition over divisible complementary factors. Ride sharing platforms compete for both riders and drivers. The value of an additional rider may depend on a firm's success in attracting divers (Rochet and Tirole, 2003). Similarly, military conflicts often involve competition for control over both airspace and the ground below it. The value of additional control over airspace may depend on the faction's level of control over the ground below it (Pirnie et al., 2005).

Success functions describes how prizes are divided among competitors as a function of resource allocations. When the success function is highly responsive to resource allocations, competitors have a stronger incentive to best respond, but adaptive models predict slower convergence to equilibrium due to nonequilibrium incentives. In contrast, the unique Nash equilibrium allocations are shown to be proportional to the prize values and do not depend on the responsiveness of the success function. To test these predictions, the

[1]See Young (2011) for more on Nash's mass action interpretation of equilibrium.

experimental design varies both the prize values and the responsiveness of the success function independently across treatment conditions.

Consistent with adaptive predictions, convergence to equilibrium was faster when success functions were less sensitive to resource allocations. Consistent with equilibrium predictions, subjects allocated more resources to compete for more valuable prizes. These results suggest that both equilibrium incentives and nonequilibrium incentives can affect strategic behavior in empirical settings. While equilibrium may provide a useful approximation of long run behavior, nonequilibrium incentives can influence the rate of convergence.

This paper contributes to the experimental literature on Blotto type contests where agents allocate limited resourced over multiple battlefields. Much of the previous literature focuses on winner-take-all battles, where a single agent receives the entirety of the prize in a given battlefield. Duffy and Matros (2017) find support for equilibrium predictions regarding differences in allocation behavior under majoritarian objective functions and linear objective functions in stochastic Blotto contests with winner-take-all battles. Chowdhury et al. (2021) find that subjects over-allocate resources to battlefields with distinctive labels in stochastic Blotto contests with winner-take-all battles. In contrast, the present paper investigates the relationship between success function responsiveness and the rate of convergence to equilibrium in Blotto contests with battles for shares of divisible complementary prizes.

This paper also contributes to a growing literature investigating the process of convergence to equilibrium. Cason et al. (2014) observe better convergence to equilibrium in evolutionarily stable games than in evolutionarily unstable games. In contrast, the games investigated by the present study are evolutionarily stable in all treatments, so evolutionary stability does not explain the observed treatment effects. Chen and Gazzale (2004) observe better convergence to equilibrium in supermodular compensation mechanisms than in non-supermodular compensation mechanisms. In contrast, the games investigated by the present study are supermodular in all treatments, so supermodularity does not explain the observed treatment effects. Stephenson (2022) observes better convergence in school choice mechanisms with high frequency feedback than in school choice mechanisms with low frequency feedback. In contrast, the present study provides feedback at the same frequency in all treatments, so feedback frequency does not explain the treatment effects identified by the present study.

The remainder of this paper is organized as follows. Section 2 presents the theory. Section 3 describes the experimental design. Section 4 states the hypotheses. Section 5 discusses the results and section 6 concludes.

2 Theory

Consider a resource allocation game where two agents simultaneously allocate competitive resources between two contests. Let $x_{ik} \in \mathbb{R}_+$ denote the quantity of competitive resources allocated by agent i to contest k. As in the Blotto contest of Borel (1921), total resource investments are sunk before agents allocate them between contests. Let $w_i \in \mathbb{R}_+$ denote the quantity of resources allocated between contests by agent i. Let X_i denote the set of allocations $x_i \in \mathbb{R}^2_+$ such that $x_{i1} + x_{i2} = w_i$. Each contest has a divisible prize. Let $v_k \in [0, 1]$ denote the relative value of contest k's prize such that $v_1 + v_2 = 1$. The success function $y_{ik}(x)$ describes agent i's share of prize k as a function of the allocation profile $x \in \mathbb{R}^{2\times 2}_+$. If $x_{1k} = x_{2k} = 0$, then prize k is evenly divided between the agents. If $x_{1k} + x_{2k} > 0$ then the success function takes the generalized Tullock (1980) form under which agent i's share of prize k is proportional to a power function of their allocation to contest k.

$$y_{ik}(x) = \frac{x^\alpha_{ik}}{x^\alpha_{ik} + x^\alpha_{jk}} \tag{1}$$

The parameter α indexes the responsiveness of the success function $y_{ik}(x)$ to the allocation levels x_{ik} and x_{jk}. If α is very large then nearly the entirety of prize k is awarded to the agent who allocates the most resources to contest k. If α is very small then prize shares are largely insensitive to allocations. Let $\pi_i(x)$ denote agent i's objective function. Prize shares are complementary inputs to the objective function. If $y_{ik}(x) = 0$ then $\pi_i(x) = 0$. If $y_i(x) \in \mathbb{R}^2_{++}$

then agent i's payoff is given by

$$\pi_i(x) = \beta \left(v_1 y_{i1}(x)^{-c} + v_2 y_{i2}(x)^{-c} \right)^{-\frac{1}{c}} \tag{2}$$

The parameter $c > 0$ indexes the level of complementarity. In the limit as c becomes large, prizes are perfect complements. Strategic interaction frequently involves competition over divisible complementary factors. For example, ride sharing platforms compete for both riders and drivers. The value of an additional rider may depend on a firm's success in attracting divers (Rochet and Tirole, 2003). Similarly, military conflicts can involve competition for control over both airspace and the ground below it. The value of additional control over the airspace in a given region may depend on a military faction's level of control over the ground below it (Pirnie et al., 2005).

This is not a zero sum game because the total payoff $\pi_1(x) + \pi_2(x)$ varies with the strategy profile x. Consider the simple case where $v_1 = v_2 = \frac{1}{2}$ and $\alpha = c = 1$. If both agents select identical resource allocations then $x_1 = x_2$ and $\pi_1(x) = \pi_1(x) = \frac{1}{2}\beta$, so the total payoff is given by $\pi_1(x) + \pi_2(x) = \beta$. In contrast if $x_1 = \left(\frac{1}{4}, \frac{3}{4}\right)$ and $x_2 = \left(\frac{3}{4}, \frac{1}{4}\right)$ then $\pi_1(x) = \pi_2(x) = \frac{3}{8}\beta$, so the total payoff is given by $\pi_1(x) + \pi_2(x) = \frac{3}{4}\beta$.

Theorem 1 says that each agent's equilibrium allocation to contest k is proportional to the value of prize k under unit complementarity. A proof of this theorem provided in the appendix. The general case with n players, m contests, and arbitrary complementarity is considered by Stephenson (2023).

Theorem 1. *If $c = 1$ then the unique Nash equilibrium satisfies $x_{ik} = w_i v_k$.*

Figure 1 illustrates the equilibrium payoff function for the resource allocation game with $c = 1$, $\beta = 28$, and $w_i = 100$. Subfigure 1a illustrates the equilibrium payoff function when the value of prize 1 is given by $v_1 = 0.2$. Subfigure 2b illustrates the equilibrium payoff function when the value of prize 1 is given by $v_1 = 0.8$. The horizontal axis indicates agent i's investment to contest 1 and the vertical axis indicates agent i's payoff. The dashed line illustrates agent i's payoff function when $\alpha = 1$. The solid line illustrates agent i's payoff function when $\alpha = 8$. The dotted line indicates agent j's

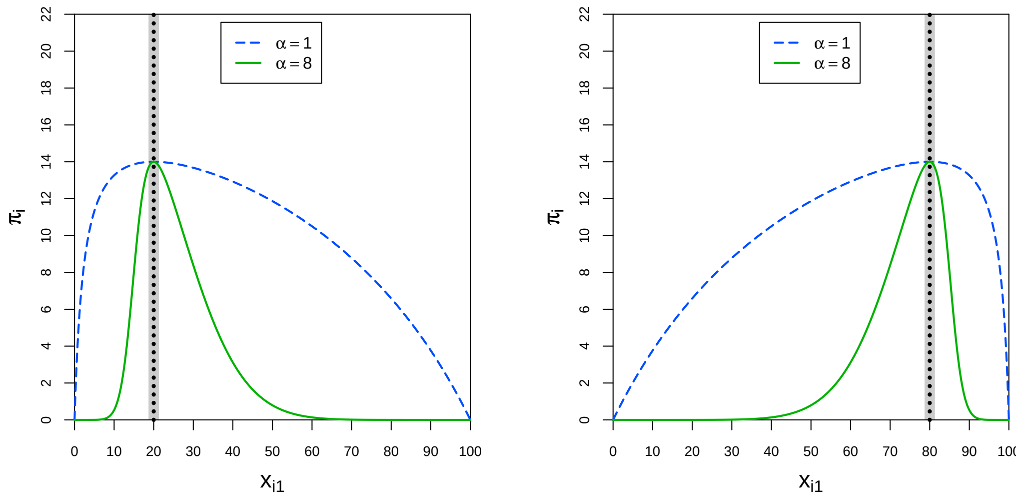

(a) $v_1 = 0.2, x_{j1} = 20$

(b) $v_1 = 0.8, x_{j1} = 80$

Figure 1: Equilibrium payoff functions for $c = 1$, $\beta = 28$, and $w_i = 100$

allocation to contest 1. The thick transparent line indicates agent i's equilibrium allocation to contest 1. The equilibrium payoff is equal to 0.5 in both cases, but payoffs are more sensitive to allocations when $\alpha = 8$.

A Nash equilibrium is said to be evolutionarily stable if small deviations from equilibrium always give the equilibrium strategy a higher payoff than than the deviating strategy (Taylor and Jonker, 1978). More formally, a symmetric Nash equilibrium (σ^*, σ^*) is said to be evolutionarily stable if, for any nonequilibrium mixed strategy $\sigma \neq \sigma^*$ and any sufficiently small $\varepsilon > 0$, $\pi_1(\sigma, \bar{\sigma}) < \pi_1(\sigma^*, \bar{\sigma})$ where $\bar{\sigma} = \varepsilon\sigma + (1 - \varepsilon)\sigma^*$ denotes a strategy that involves utilizing the the nonequilibrium strategy σ with probability ε and utilizing the equilibrium strategy σ^* with probability $1 - \varepsilon$. Intuitively, such equilibria are stable because small deviations from equilibrium never incentivize equilibrium players to adopt the deviating strategy. As shown in the proof of theorem 1, the objective function π_i is strictly quasiconcave in x_i, so the Nash equilibrium x^* is always strict and the equilibrium strategy is always evolutionarily stable.

A symmetric two player game with a one dimensional strategy space is said

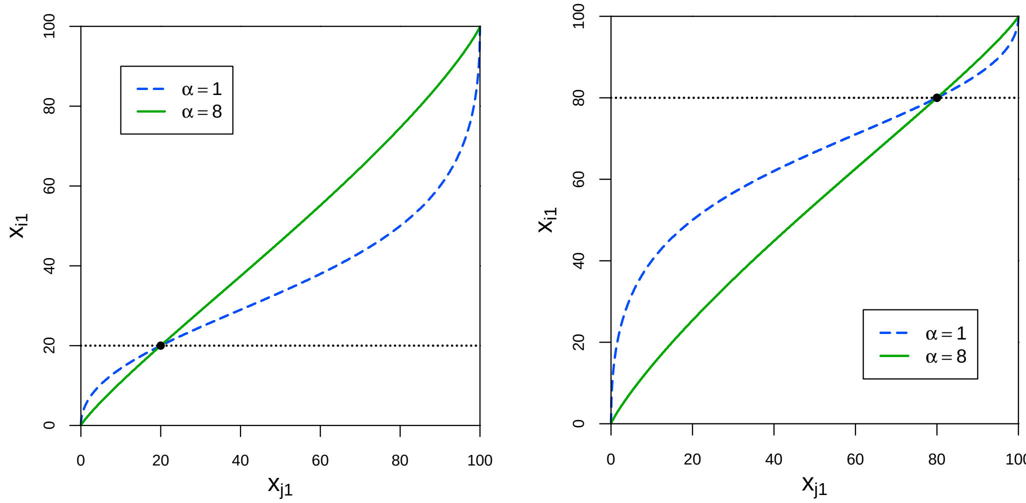

(a) $v_1 = 0.2$ (b) $v_1 = 0.8$

Figure 2: Best response functions for $c = 1$ and $w_i = 100$

to be supermodular if the marginal payoff to one player from increasing their strategy is increasing in the other player's strategy. The resource allocation game described above is symmetric and the strategy space is given by the one dimensional unit simplex $\triangle_1$. If $c = 1$ then differentiating player i's marginal benefit from allocation to contest 1 with respect to agent j's allocation to contest 1 yields

$$\frac{\partial^2 \pi_i}{\partial x_{i1} \partial x_{j1}} = \alpha \beta \pi_i \left(x\right)^2 \left[\frac{v_1}{x_{i1}^2} + \frac{v_2}{x_{i2}^2} \right] \tag{3}$$

This expression is strictly positive for all $x_i \in \mathbb{R}^2_{++}$, so player i's marginal benefit from allocating resources to prize 1 is increasing in agent j's allocation to prize 1. Hence the resource allocation game with $c = 1$ is supermodular for all $\alpha > 0$.

Figure 2 illustrates the best response correspondence. The horizontal axis indicates agent j's allocation to contest 1 and the vertical axis indicates agent i's payoff maximizing allocation to contest 1. Subfigure 2a illustrates agent i's best response function when the value of prize 1 is given by $v_1 =$

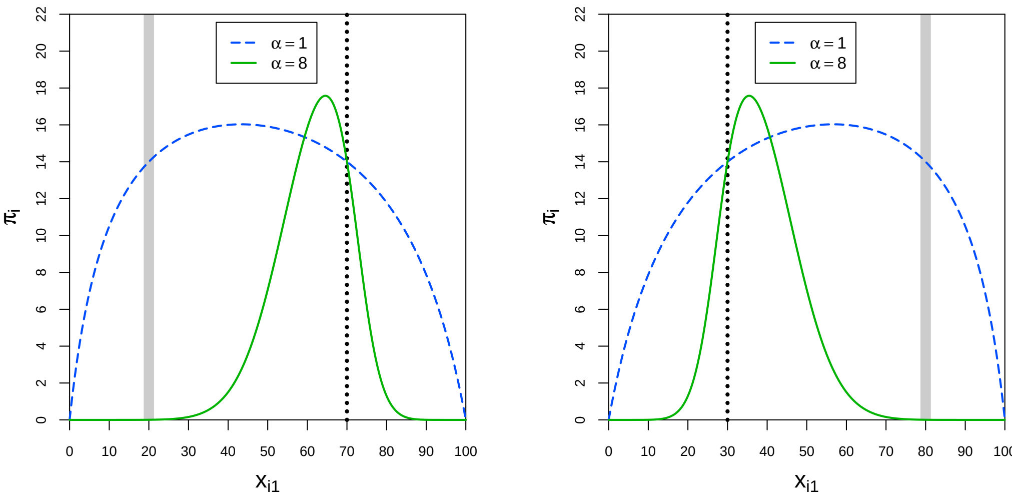

(a) $v_1 = 0.2,\ x_{j1} = 70$ (b) $v_1 = 0.8,\ x_{j1} = 30$

Figure 3: Nonequilibrium payoff functions for $c = 1$, $\beta = 28$, and $w_i = 100$

0.2. Subfigure 2b illustrates agent i's best response function when the value of prize 1 is given by $v_1 = 0.8$. The dashed line illustrates agent i's best response when $\alpha = 1$. The solid line indicates agent i's best response when $\alpha = 0.8$. The dotted line indicates the equilibrium allocation to prize 1. Out of equilibrium, agent i's best response is always closer to equilibrium when $\alpha = 1$ than $\alpha = 8$.

Figure 3 illustrates nonequilibrium payoff functions in resource allocation games with unit endowments. Subfigure 3a illustrates agent i's payoff function when the value of prize 1 is given by $v_1 = 0.2$ and agent j's allocation to contest 1 is given by $x_{j1} = 0.7$. Subfigure 3b illustrates agent i's payoff function when the value of prize 1 is given by $v_1 = 0.8$ and agent j's allocation to contest 1 is given by $x_{j1} = 0.3$. The horizontal axis indicates agent i's investment in prize 1 and the vertical axis indicates agent i's payoff. The dashed line illustrates agent i's payoff function when $\alpha = 1$. The solid line illustrates agent i's payoff function when $\alpha = 8$. The dotted vertical line indicates the agent j's allocation to contest 1. The thick transparent vertical line indicates the equilibrium allocation to contest 1. Agent i's payoff maximizing allocation is closer to equilibrium when $\alpha = 1$ than $\alpha = 8$.

2.1 Quantal Response Equilibrium

Quantal response equilibrium describes boundedly rational agents who are more likely to select strategies that yield higher payoffs (McKelvey and Palfrey, 1995). In contrast, Nash equilibrium describes agents who always best respond. Let $\bar{\pi}_i(x_i|\sigma_j)$ denote agent i's expected payoff from the pure strategy x_i given agent j's mixed strategy σ_j.

$$\bar{\pi}_i(x_i|\sigma_j) = \int \pi_i(x_i, x_j)\,\sigma_j(x_j)\,dx_j \tag{4}$$

Agents i's logit quantal response to the mixed strategy σ_j is given by the probability density function $\ell_i(x_i|\sigma_j)$.

$$\ell_i(x_i|\sigma_j) = \frac{\exp\lambda\bar{\pi}_i(x_i|\sigma_j)}{\int \exp\lambda\bar{\pi}_i(y_i|\sigma_j)\,dy_i} \tag{5}$$

The logit parameter λ describes the sensitivity of the logit quantal response to payoff differences. In the limit as $\lambda \to \infty$, the logit quantal response converges to the best response. In the limit as $\lambda \to 0$, the logit quantal response converges to uniformly random strategy selection. A logit quantal response equilibrium is a fixed point σ^* of the logit quantal response such that $\sigma_i^*(x_i) = \ell_i\left(x_i|\sigma_j^*\right)$. As shown in figure 1, the responsiveness parameter α affect the shape of the payoff function without affecting the equilibrium best response, so changes in α produce changes in the logit quantal response equilibrium without changing the Nash equilibrium predictions.

Figure 4 illustrates logit quantal response equilibria of the resource allocation game with $c = 1$, $\beta = 28$, $w_i = 100$, and $v_1 = 0.8$. The horizontal axis of each graph depicts the allocation to contest 1 and the vertical axis depicts the logit quantal response equilibrium density. The thick transparent line indicates the Nash equilibrium allocation to contest 1. The dotted line indicates the average logit quantal response equilibrium allocation to contest 1. The shape of the logit quantal response equilibrium distribution depends on both α and λ. Higher values of α make payoffs more sensitive to allocations and higher values of λ make allocations more sensitive to payoffs.

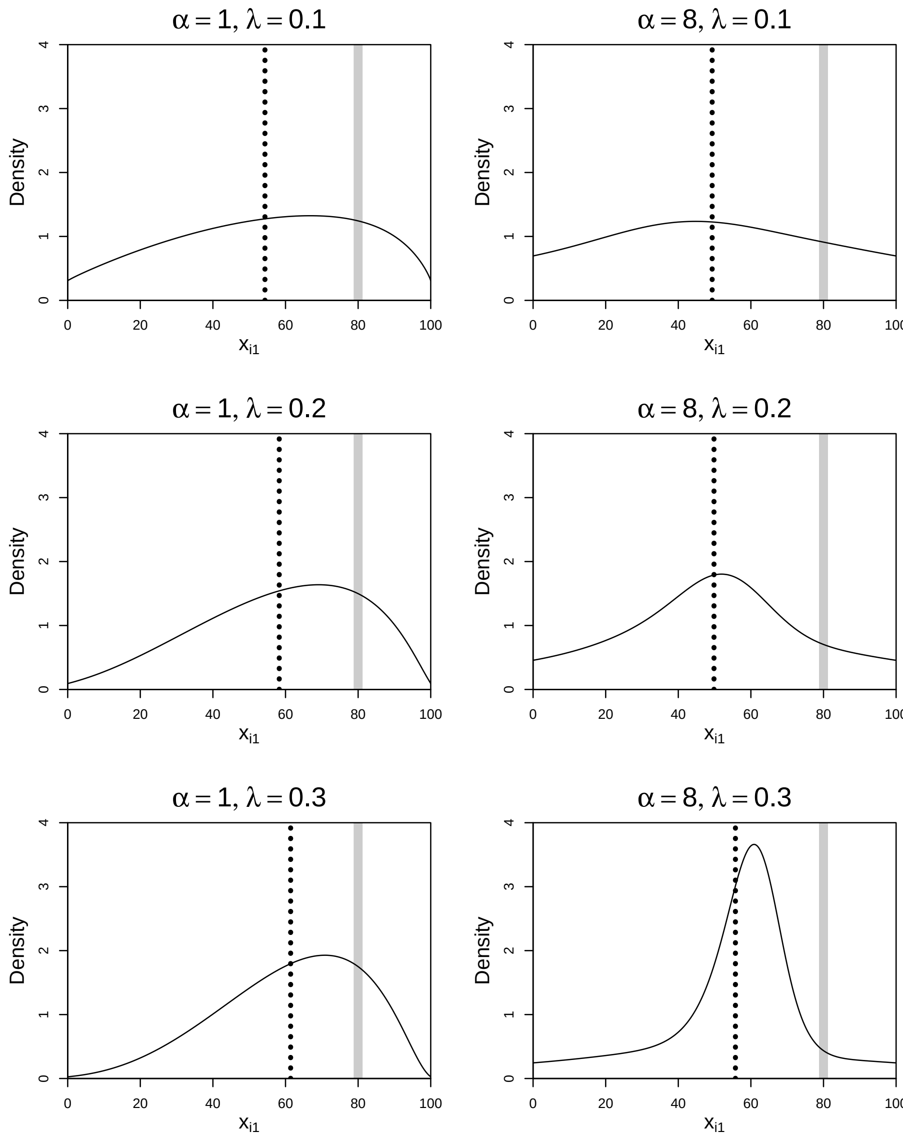

Figure 4: Logit quantal response equilibria for $c = 1$, $\beta = 28$, $w_i = 100$, and $v_1 = 0.8$

2.2 Adaptive Models

The existence of a unique Nash equilibrium implies that any finite repetition of the resource allocation game has a unique subgame perfect Nash equilibrium repeating the stage game equilibrium in every period. In contrast, Nash's "mass action" interpretation of equilibrium considers boundedly rational agents who "accumulate empirical information on the relative advantages of the various pure strategies at their disposal" (Nash, 1950). In this case, "it is unnecessary to assume that participants have full knowledge of the total structure of the game, or the ability and inclination to go though any complex reasoning process" (Nash, 1950). If agents exhibit this kind of bounded rationality, Nash equilibrium may provide a more accurate description of long run behavior than short run behavior.

Adaptive models formalize this intuition by explicitly describing how agents change their behavior over time. The logit dynamic describes agents who quantal respond to the action taken by their opponent in the previous period (Fudenberg and Levine, 1998). Let $\sigma_{it}^{L}(x_i)$ denote the probability density for agent i's allocation $x_i(t)$ in stage $t > 1$ under the logit dynamic.

$$\sigma_{it}^{L}(x_i) = \frac{\exp \lambda \pi_i (x_i, x_j(t-1))}{\int \exp \lambda \pi_i (y_i, x_j(t-1))\, dy_i} \tag{6}$$

The noisy best response dynamic describes agents who are sensitive to both the magnitude of payoff differences and the location of the best response. Let $\sigma_{it}^{N}(x_i)$ denote the probability density for agent i's allocation $x_i(t)$ in stage $t > 1$ under the noisy best response dynamic.

$$\sigma_{it}^{N}(x_i) = \frac{\exp u_{it}(x_i, x_j(t-1))}{\int \exp u_{it}(y_i, x_j(t-1))\, dy_i} \tag{7}$$

$$u_{it}(x_i, x_j) = \lambda \pi_i (x_i, x_j) - \eta \left| x_i - x_i^*(x_j) \right| - \gamma \left| x_i - x_i(t-1) \right| \tag{8}$$

$$x_i^*(x_j) = \underset{x_i \in X_i}{\operatorname{argmax}}\, \pi_i(x_i, x_j) \tag{9}$$

The parameter λ describes agent i's sensitivity to the magnitude of payoff differences. The parameter η describes agent i's sensitivity to the location

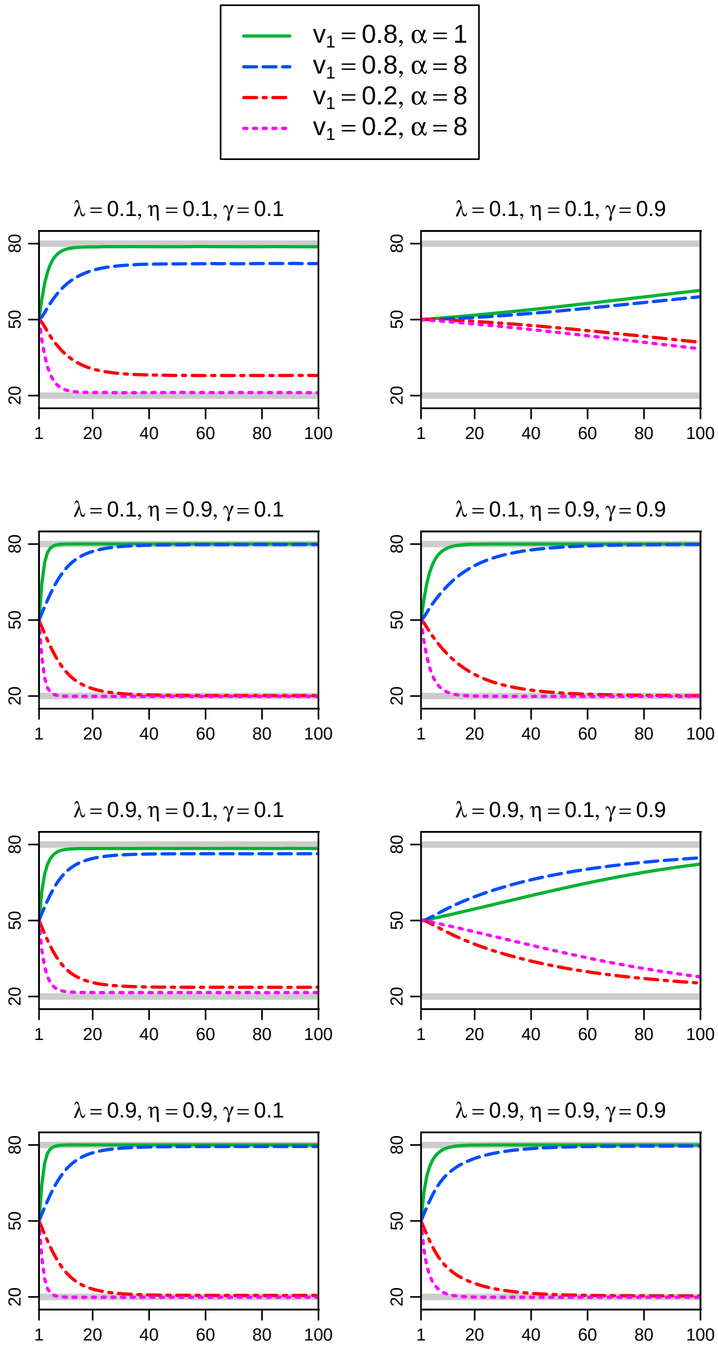

Figure 5: Average path predictions for the noisy best response dynamic.

of the best response. The parameter γ describes the strength of behavioral inertia. Behavioral inertia is the tendency for agents to continue doing what they did in the past (Norman, 2009; Liu and Riyanto, 2017). In the limit as $\lambda \to 0$, $\eta \to 0$, and $\gamma \to 0$, agent i's allocation is uniformly distributed. In the limit as $\lambda \to \infty$ or $\eta \to \infty$, agent i's allocation coincides with the best response. In the limit as $\gamma \to \infty$, agent i permanently maintains their original allocation.

Figure 5 illustrates the predictions of of the noisy best response dynamic. The horizontal axis of each graph indicates the period. The vertical axis of each graph illustrates the average allocation to contest 1. The solid lines indicate the predictions for $v_1 = 0.8$ and $\alpha = 1$. The dashed lines illustrates the predictions for $v_1 = 0.8$ and $\alpha = 8$. The dot-dashed lines indicate the predictions for $v_1 = 0.2$ and $\alpha = 8$. The dotted lines illustrates the predictions for $v_1 = 0.2$ and $\alpha = 1$.

3 Experimental Design

The experiment has a 2×2 factorial design with a total of 4 treatment conditions as illustrated by table 1. In the low responsiveness treatment, the responsiveness of the success function was given by $\alpha = 1$. In the high responsiveness treatment, the responsiveness of the success function was given by $\alpha = 8$. Section 2 provides a detailed discussion of the responsiveness parameter α. In the first valuation treatment, prize values were given by $v = (0.8, 0.2)$, so the unique Nash equilibrium predicts that agents will allocate 80% of their resources to contest 1. In the second valuation treatment, prize values were given by $v = (0.2, 0.8)$, so the unique Nash equilibrium predicts that agents will allocate 20% of their resources to contest 1.

Each experimental session implemented one of the four treatment conditions. A total of 8 experimental sessions were conducted, 2 for each of the 4 treatment conditions. Each session had 20 subjects for a total of 160 experimental subjects. At the beginning of each session, subjects were randomly matched into pairs which remained fixed for the entire session. Each experimental session consisted of 100 periods. The first period lasted for 1 minute. Each

	Low Responsiveness	High Responsiveness
First Valuation	$\alpha = 1,\, v_1 = 0.8$	$\alpha = 8,\, v_1 = 0.8$
Second Valuation	$\alpha = 1,\, v_1 = 0.2$	$\alpha = 8,\, v_1 = 0.2$

Table 1: Experimental Design

of the subsequent periods lasted for 5 seconds.

Each period implemented the resource allocation game described in section 2. During a period, each subject allocated 100 tokens between two contests. At the end of the period, the payoff to each subject depended on both how they allocated their resources and how their opponent allocated resources. Consistent with Benndorf et al. (2016), Leng et al. (2018), and Cason et al. (2021), subjects were shown the allocation selected by each group member, the payoff earned by each group member, and the payoffs they could have earned by selecting other allocations at the end of each period.

Figure 6 depicts the experimental interface. In the central graph, the horizontal axis indicates the number of tokens invested in contest 1 and the vertical axis indicates the subject's payoff. The black line indicates the number of tokens the subject chose to invest in contest 1 during the previous period. The green line illustrates the payoffs the subject could have earned in the previous period if they had selected other allocations. The blue line indicates the number of tokens the subject has currently selected to invest in contest 1 during the current period. Numerical information about allocations and payoffs are shown below the graph. In each period, payoffs were determined by the objective function given by equation (2) multiplied by 28. At the end of a session, subjects received their average payoff over all 100 periods plus a $7 participation bonus. Average earnings in the experiment were $19.92 per subject.

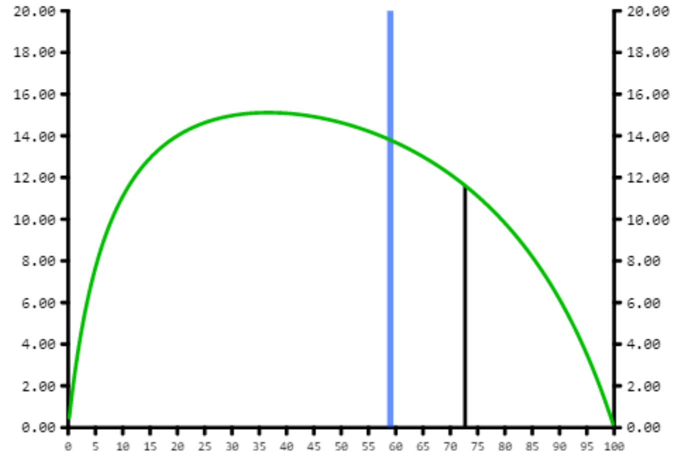

Figure 6: Experimental Interface

4 Hypotheses

The unique Nash equilibrium of the game described in section 2 has players allocating resources to each contest in proportion to the value of it's prize. In the first valuation treatment, Nash equilibrium predicts that subjects will invest 80% of their resources in contest 1. In the second valuation treatment, it predicts that subjects will invest 20% of their resources in contest 1.

Hypothesis 1. *More resources will be allocated to contest 1 in the first valuation treatment than the second valuation treatment.*

In equilibrium, agents have a stronger incentive to best respond when the contest success function is more responsive to investment levels, but adaptive models often predict faster convergence to equilibrium under less responsive success functions. As discussed in section 2.2, nonequilibrium best responses are consistently farther from equilibrium under more responsive success functions.

Hypothesis 2. *Resource allocations will approach equilibrium more quickly in the low responsiveness treatment than the high responsiveness treatment.*

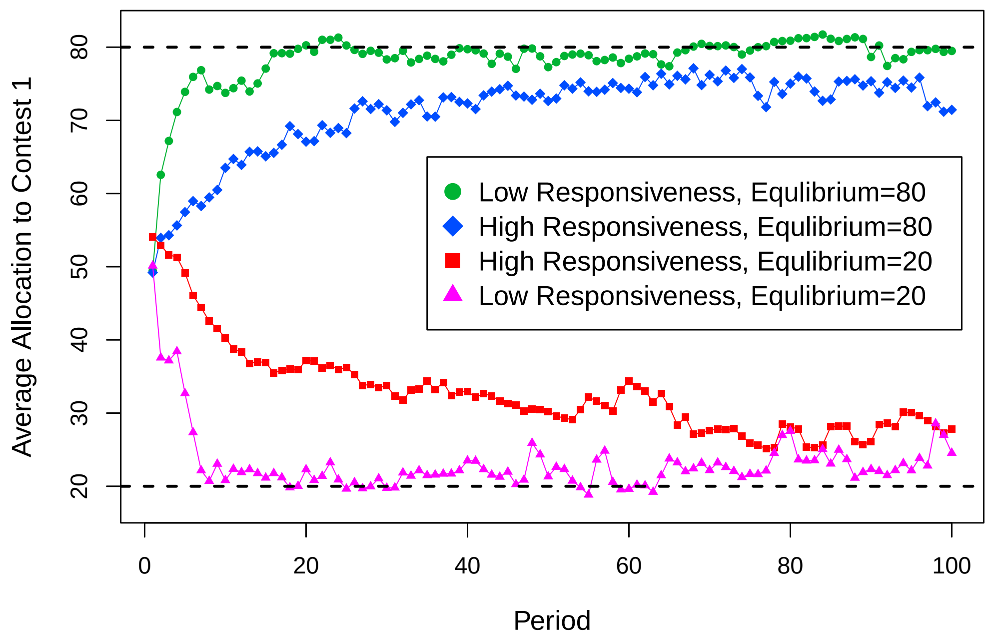

Figure 7: Average allocations by period

5 Results

Figure 7 illustrates average allocations by period under each of the four treatment conditions. The vertical axis indicates the average allocation to contest 1. The horizontal axis indicates the period. In the first valuation treatment, the unique Nash equilibrium predicts that agents will allocate 80% of their resources to contest 1. In the second valuation treatment, the unique Nash equilibrium predicts that agents will allocate 20% of their resources to contest 1. Each subject allocated a total of 100 units between the two contests. The average allocation to contest 1 under the first valuation treatment was 74.72 while the average allocation to contest 1 in the second valuation treatment was 27.9.

Consistent with hypothesis 1, subjects allocated significantly more resources to contest 1 in the first valuation treatment than the second valuation treatment. Both a t-test and a non-parametric Wilcoxon rank-sum test find this difference to be statistically significant at the 1% level, as reported in table

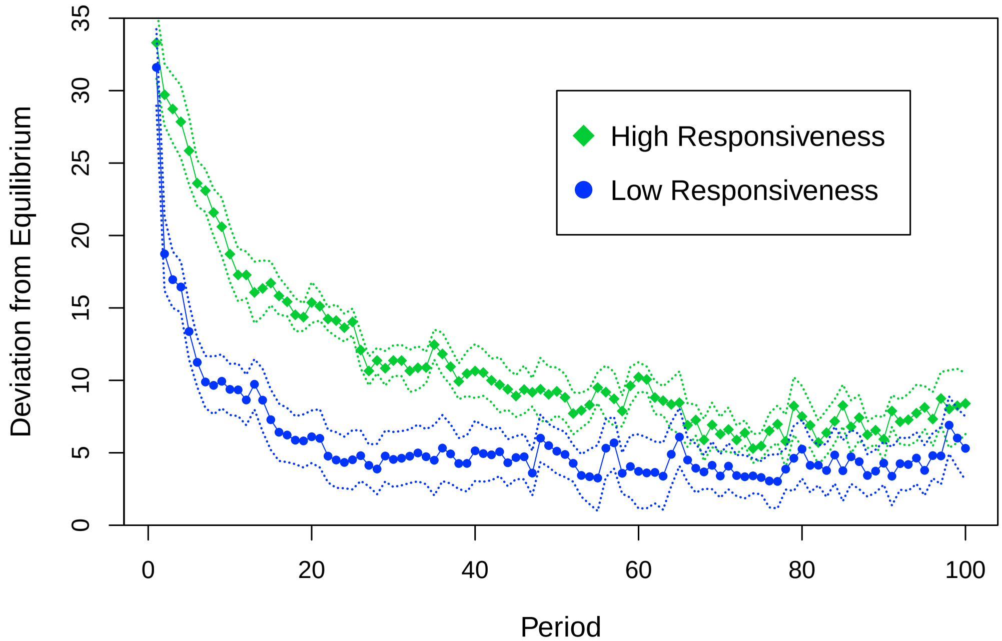

Figure 8: Deviation from equilibrium by period

2. The average allocation selected by a fixed matching pair over an entire experimental session was treated as a single observation. A total of 8 sessions were conducted. There were 4 sessions per valuation treatment and 10 fixed matching pairs per session, yielding a total of 40 observations with 20 observations per valuation treatment.

Result 1. *Significantly more resources were allocated to contest 1 in the first valuation treatment than the second valuation treatment.*

Figure 8 illustrates the average deviation from equilibrium by period in each of the two responsiveness treatments. Deviation from equilibrium is measured as the absolute difference between the observed allocation to contest 1 and the equilibrium allocation to contest 1. The vertical axis indicates the average deviation from equilibrium. The horizontal axis indicates the period. In the high responsiveness treatment, the responsiveness of the contest success function was given by $\alpha = 8$. In the low responsiveness treatment, the responsiveness of the contest success function was given by $\alpha = 1$. Dotted

			p-value	
Responsiveness (α)	1	8	rank-sum	t-test
Deviation from Equilibrium	5.73	11.21	< 0.0001	< 0.0001
Equilibrium Allocation 1	20	80	rank-sum	t-test
Average Allocation 1	27.9	74.72	< 0.0001	< 0.0001

Table 2: Pair Level Hypothesis Tests

lines illustrate standard errors for the average deviation from equilibrium under each responsiveness treatment in each period.

Consistent with hypothesis 2, resource allocations were significantly closer to equilibrium predictions in the low responsiveness treatment than the high responsiveness treatment. The average deviation from equilibrium in the low responsiveness treatment was 5.73 while the average deviation from equilibrium in the high responsiveness treatment was 11.21. Both a t-test and a non-parametric Wilcoxon rank-sum test find this difference to be statistically significant at the 1% level, as reported in table 2. We treat the average allocation selected by a fixed matching pair over an entire experimental session as a single observation, yielding a total of 40 observations with 20 observations per responsiveness treatment.

Result 2. *Resource allocations were significantly closer to equilibrium predictions in the low responsiveness treatment than the high responsiveness treatment.*

Adaptive models predict faster convergence to equilibrium in the low responsiveness treatment than the high responsiveness treatment, as discussed in section 2.2. These models describe a gradual process of behavioral change as agents accumulate experience, but they make no predictions about the initial resource allocations selected by subjects during the first period. In contrast, cognitive hierarchy and level-k models can characterize the initial behavior of boundedly rational agents who exhibit a limited depth of introspective

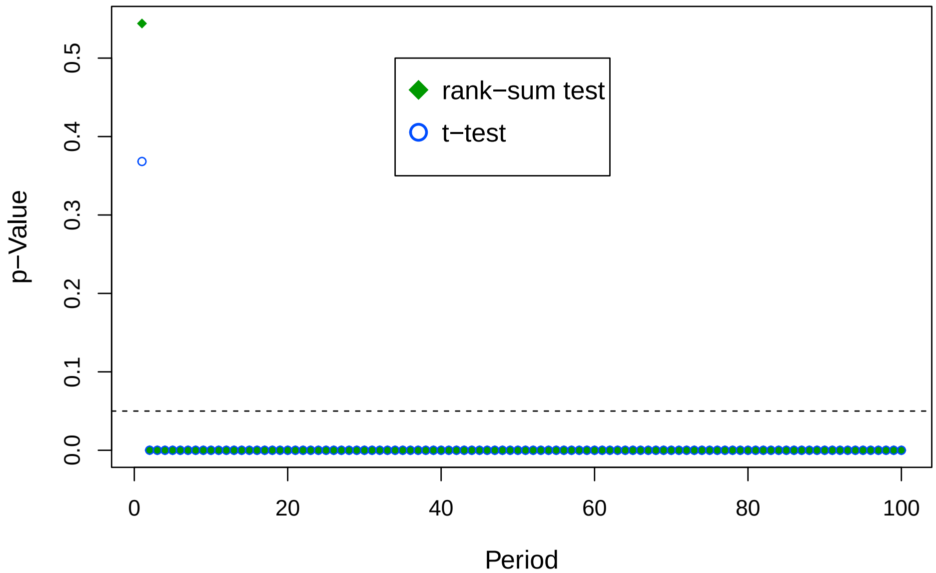

Figure 9: p-values for differences in allocations across valuation treatments

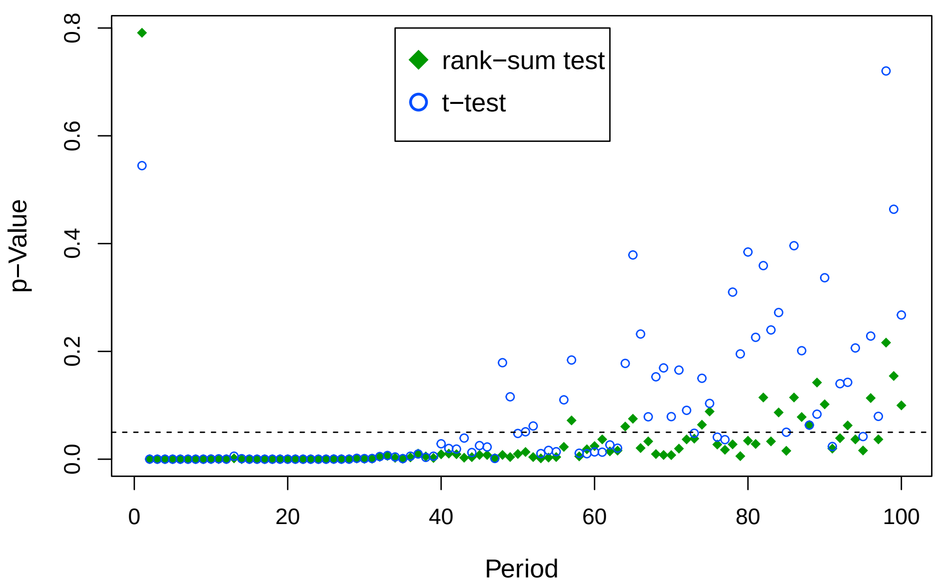

Figure 10: p-values for differences in deviation from equilibrium across responsiveness treatments

reasoning (Nagel, 1995; Camerer et al., 2004). Accordingly, cognitive hierarchy models can predict initial behavior that is closer to equilibrium under less responsive contest success functions. No significant differences are found in initial behavior across treatment conditions, suggesting that the observed treatments effects are driven by differences in the way behavior changes over time, rather than differences in the distribution of initial behavior.

Figures 9 and 10 illustrate p-values for hypothesis tests using data from only a single period. Figure 9 illustrates tests for differences in allocation levels across valuation treatments and figure 10 illustrates tests for differences in distance from equilibrium across responsiveness treatments. In both figures, the dashed line indicates the conventional 5% significance level. In the first period, there is no significant difference in resource allocations or deviations from equilibrium, suggesting that the observed treatment effects are not driven by introspective reasoning prior to initial play. In all subsequent periods, both tests find significantly higher allocations to contest 1 in the first valuation treatment. In periods 2-40, both tests consistently find allocations to be significantly farther from equilibrium in the high responsiveness treatment. In periods 80-100, neither test consistently finds significant differences in deviation from equilibrium across responsiveness treatments. Neither initial behavior in the first period nor long run behavior in the last few periods is significantly different across responsiveness treatments, suggesting that result 2 is driven by differences in the rate of convergence during intermediate periods.

Parameters for the adaptive model described in section 2.2 are estimated at the subject level by maximum likelihood. Table 3 provides treatment level averages and hypothesis tests. Hypothesis tests treat each fixed matching group as a single observation for a total of 80 observations with 8 sessions and 10 fixed matching groups per session. Table 3a provides the average value of each parameter estimate over all experimental subjects. Standard errors are provided in parentheses. Signed-ranks tests find that subjects exhibited significant sensitivity to payoff differences, significant sensitivity to the location of the best response, and significant behavioral inertia at the 1% level. Table 3b provides average parameter values by treatment. Rank-sum tests find no significant effect from responsiveness of the success function

Parameter	Estimate	Signed-Rank Test p-value
Payoff Sensitivity (λ)	4.201 (1.451)	0.009
Best Response Sensitivity (η)	0.770 (0.130)	<0.001
Behavioral Inertia (γ)	0.949 (0.157)	<0.001

(a) Average Parameter Estimates

	Responsiveness		Rank-Sum Test
Parameter	$\alpha = 1$	$\alpha = 8$	p-value
Payoff Sensitivity (λ)	8.185 (2.771)	0.217 (0.194)	0.121
Best Response Sensitivity (η)	0.536 (0.119)	1.005 (0.228)	0.72
Behavioral Inertia (γ)	1.023 (0.223)	0.875 (0.223)	0.034

(b) Parameter Estimates by Treatment

	$\alpha = 1$			$\alpha = 8$		
	λ	η	γ	λ	η	γ
Payoff Sensitivity (λ)	1.00	0.59	0.87	1.00	-0.12	0.29
Best Response Sensitivity (η)	0.59	1.00	0.84	-0.12	1.00	0.86
Behavioral Inertia (γ)	0.87	0.84	1.00	0.29	0.86	1.00

(c) Correlation Between Parameter Estimates

Table 3: Parameter Estimates at the Subject Level

on the level of payoff sensitivity or the level of sensitivity to the location of the best response. In contrast, a rank-sum test finds that subjects exhibited significantly higher levels of behavioral inertia when $\alpha = 1$ than when $\alpha = 8$ at the 5% level. Stronger behavioral inertia may have resulted from the relatively weaker incentives to best respond under less responsive success functions, as illustrated in Figure 3.

Table 3c provides correlations between subject level parameter estimates. Payoff sensitivity, best response sensitivity, and behavioral inertia were all positively correlated in the low responsiveness treatment where $\alpha = 1$. Positive correlation between sensitivity levels and behavioral inertia may indicate that subjects face a tradeoff between speed and precision of their allocation adjustments. Subjects who make careful adjustments might tend to make large adjustments less frequently. Payoff sensitivity and best response sensitivity were negatively correlated in the high responsiveness treatment where $\alpha = 8$, suggesting that subjects with limited attention may face a tradeoff between paying attention to the location of the best response and paying attention to the magnitude of payoff differences in settings where outcomes are highly sensitive to actions.

6 Conclusions

To predict the impact of a hypothetical potential policy intervention, economists often consider how it would affect equilibrium predictions. Yet the transition from a pre-intervention equilibrium to a post-intervention equilibrium may be gradual if economic agents are boundedly rational, consistent with the mass action interpretation of equilibrium described by Nash (1950). Adaptive models formalize this idea by explicitly describing how boundedly rational agents adjust their behavior over time.

This study experimentally tests both equilibrium predictions and adaptive predictions in settings where agents allocate limited resources between contests with divisible complementary prizes. The unique Nash equilibrium resource allocations are shown to be proportional to prize values and do not depend on the responsiveness of the success function. In contrast, adaptive models predict slower convergence to equilibrium under more responsive success functions due to nonequilibrium incentives.

To test these predictions, the experimental design varies both the prize values and the responsiveness of the success function independently across treatment conditions. Consistent with equilibrium predictions, subjects allocated more resources to compete for more valuable prizes. Consistent with adaptive model predictions, convergence to equilibrium was faster when success functions were less responsive to resource allocations. These results suggest that both equilibrium incentives and nonequilibrium incentives contain important information for policy makers. While equilibrium may provide a useful tool for characterizing long run behavior, nonequilibrium incentives can affect the rate of convergence.

The present study investigates one particular class of strategic environments, but additional research is needed to better understand the factors that determine the rate of convergence in a wider variety of settings before strong conclusions can be drawn about the generality of the present results. The experimental design of the present study varies both the prize values and the responsiveness of the success function, but it does not vary the number of competitors, the number of prizes, or the elasticity of substitution between prizes. Future research should investigate how these other factors effect the rate of convergence. The present study also finds that the responsiveness of the success function had a significant effect on the strength of behavioral inertia. Additional research is needed to better understand the underlying factors determining the strength of behavioral inertia in strategic settings.

A Proofs

Proof of Theorem 1. Suppose agent i allocates zero resources to contest k such that $x_{ik} = 0$. Let $\hat{x}_i$ such that $\hat{x}_{ik} = \varepsilon \in (0, 1)$ and $\hat{x}_{ib} = 1 - \varepsilon$. If $x_{ik} = 0$ and $x_{jk} \neq 0$ then $\pi_i(x) = 0 < \pi_i(\hat{x}_i, x_j)$. If $x_{ik} = x_{jk} = 0$ then taking the limit as $\varepsilon \to 0$ obtains.

$$\lim_{\varepsilon \to 0} \pi_i\left(\hat{x}_i, x_j\right) = \frac{\beta}{v_k + 2v_b} > \frac{\beta}{2} = \pi_i\left(x\right) \tag{10}$$

Hence $x_{ik} > 0$ in every Nash equilibrium. Let $g_i\left(x\right) = -\frac{\beta}{\pi_i(x)}$ so differentiating g_i with respect to x_{ik} yields

$$\frac{\partial g_i}{\partial x_{ik}} = \frac{\alpha v_k\left[1 - y_{ik}\left(x\right)\right]}{y_{ik}\left(x\right) x_{ik}} \tag{11}$$

Since $y_{ik}\left(x\right)$ is increasing in x_{ik}, the numerator of (11) is decreasing in x_{ik} and the denominator is increasing in x_{ik} so $\frac{\partial g_i^2}{\partial x_{ik}^2} < 0$. If $b \neq k$ then (11) is constant in x_{ib} so $\frac{\partial g_i^2}{\partial x_{ib}^2} = 0$. Hence g_i is strictly concave in x_i and π_i is strictly quasiconcave in x_i. The first order conditions on x_i for the maximization of π_i state that $\frac{\partial \pi_i}{\partial x_{ik}} = \frac{\partial \pi_i}{\partial x_{ib}}$ so we have

$$\frac{v_k\left[1 - y_{ik}\left(x\right)\right]}{y_{ik}\left(x\right) x_{ik}} = \frac{v_b\left[1 - y_{ib}\left(x\right)\right]}{y_{ib}\left(x\right) x_{ib}} \tag{12}$$

$$\frac{v_k y_{jk}\left(x\right)}{y_{ik}\left(x\right) x_{ik}} = \frac{v_b y_{jb}\left(x\right)}{y_{ib}\left(x\right) x_{ib}} \tag{13}$$

$$\frac{v_k y_{jk}\left(x\right) y_{ib}\left(x\right)}{v_b y_{jb}\left(x\right) y_{ik}\left(x\right)} = \frac{x_{ik}}{x_{ib}} \tag{14}$$

Since π_i is strictly quasiconcave in x_i and $x_{ik} > 0$ in every equilibrium, the first order conditions are necessary and sufficient for equilibrium so

$$\frac{x_{ik}}{x_{ib}} = \frac{x_{jk}}{x_{jb}} \tag{15}$$

$$\frac{x_{ik}}{x_{jk}} = \frac{x_{ib}}{x_{jb}} \tag{16}$$

$$\frac{x_{ik}^{\alpha}}{x_{ik}^{\alpha} + x_{jk}^{\alpha}} = \frac{x_{ib}^{\alpha}}{x_{ib}^{\alpha} + x_{jb}^{\alpha}} \tag{17}$$

$$y_{ik}\left(x\right) = y_{ib}\left(x\right) \tag{18}$$

Hence there exists $\bar{y}_i\left(x\right) = y_{ik}\left(x\right) = y_{ib}\left(x\right)$. Substituting this into equation

(14) yields

$$\frac{v_k \bar{y}_j(x)\, \bar{y}_i(x)}{v_b \bar{y}_j(x)\, \bar{y}_i(x)} = \frac{x_{ik}}{x_{ib}} \tag{19}$$

$$\frac{v_k}{v_b} = \frac{x_{ik}}{x_{ib}} \tag{20}$$

Now since $x_{i1} + x_{i2} = w_i$ we have $x_{ik} = w_i v_k$. □

References

Volker Benndorf, Ismael Martinez-Martinez, and Hans-Theo Normann. Equilibrium selection with coupled populations in hawk–dove games: Theory and experiment in continuous time. *Journal of Economic Theory*, 165: 472–486, 2016.

Emile Borel. La théorie du jeu et les équations intégralesa noyau symétrique. *Comptes rendus de l'Académie des Sciences*, 173(1304-1308):58, 1921. English translation by Leonard J. Savage in Econometrica 21(1):97-100, 1953.

Colin F Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004.

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

Timothy N Cason, Daniel Friedman, and Ed Hopkins. An experimental investigation of price dispersion and cycles. *Journal of Political Economy*, 129(3):789–841, 2021.

Yan Chen and Robert Gazzale. When does learning in games generate convergence to nash equilibria? the role of supermodularity in an experimental setting. *American Economic Review*, 94(5):1505–1535, 2004.

Subhasish M Chowdhury, Dan Kovenock, David Rojo Arjona, and Nathaniel T Wilcox. Focality and asymmetry in multi-battle contests. *The Economic Journal*, 131(636):1593–1619, 2021.

John Duffy and Alexander Matros. Stochastic asymmetric blotto games: An experimental study. *Journal of Economic Behavior & Organization*, 139: 88–105, 2017.

Drew Fudenberg and David K Levine. *The theory of learning in games*, volume 2. MIT press, 1998.

John Maynard Keynes. *A tract on monetary reform*. Cosimo Classics, 1923.

Ailin Leng, Lana Friesen, Kenan Kalayci, and Priscilla Man. A minimum effort coordination game experiment in continuous time. *Experimental Economics*, 21:549–572, 2018.

Jia Liu and Yohanes E Riyanto. The limit to behavioral inertia and the power of default in voluntary contribution games. *Social choice and welfare*, 48 (4):815–835, 2017.

Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.

Rosemarie Nagel. Unraveling in guessing games: An experimental study. *The American economic review*, 85(5):1313–1326, 1995.

John Nash. *Non-cooperative games*. PhD thesis, Princeton University, 1950.

Thomas WL Norman. Rapid evolution under inertia. *Games and Economic Behavior*, 66(2):865–879, 2009.

Bruce R Pirnie, Alan Vick, Adam Grissom, Karl P Mueller, and David T Orletsky. Beyond close air support. forging a new air-ground partnership. Technical report, Rand Corporation, Santa Monica, CA, 2005.

Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the european economic association*, 1(4):990–1029, 2003.

Daniel Stephenson. Assignment feedback in school choice mechanisms. *Experimental Economics*, 25(5):1467–1491, 2022.

Daniel Stephenson. Multi-battle contests over complementary battlefields. 2023. Working Paper.

Peter D Taylor and Leo B Jonker. Evolutionarily stable strategies and game dynamics. *Mathematical biosciences*, 40(1-2):145–156, 1978.

G. Tullock. Efficient rent-seeking. In J.M. Buchanan, R.D. Tollison, and G. Tullock, editors, *Toward a Theory of the Rent Seeking Society*, pages 97–112. Texas A & M University Press, 1980.

H Peyton Young. Commentary: John nash and evolutionary game theory. *Games and Economic Behavior*, 71(1):12–13, 2011.