

NBER WORKING PAPER SERIES

STARTING SMALL: ENDOGENOUS STAKES AND RATIONAL COOPERATION

James Andreoni  
Michael A. Kuhn  
Larry Samuelson

Working Paper 21934  
<http://www.nber.org/papers/w21934>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
January 2016

Financial support from the National Science Foundation (SES-1427355 to Andreoni, and SES-1357780 to Samuelson) is gratefully acknowledged, and advice from Doug Miller and David Eil are appreciated. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by James Andreoni, Michael A. Kuhn, and Larry Samuelson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Starting Small: Endogenous Stakes and Rational Cooperation  
James Andreoni, Michael A. Kuhn, and Larry Samuelson  
NBER Working Paper No. 21934  
January 2016  
JEL No. C92,D64,Z13

**ABSTRACT**

We report experimental results for a twice-played prisoners' dilemma in which the players can choose the allocation of the stakes across the two periods. Our point of departure is the assumption that some (but not all) people are principled to "do the right thing," or cooperate, as long as their opponent is sufficiently likely to do so. The presence of such types can be exploited to enhance cooperation by structuring the twice-played prisoners' dilemma to "start small," so that the second-stage stakes are larger (but not too much larger) than the first-stage stakes. We compare conditions where the allocation of stakes is chosen exogenously to conditions where it is chosen by the players themselves. We show that players are able to find and choose the payoff maximizing strategy of starting small in a twice-played prisoners' dilemma, and that the salutary payoff effects of doing so are larger than those that arise when the same allocation is exogenously chosen.

James Andreoni  
Department of Economics  
University of California, San Diego  
9500 Gilman Drive  
La Jolla, CA 92093-0508  
and NBER  
andreoni@ucsd.edu

Larry Samuelson  
Department of Economics  
Yale University  
Box 208281  
New Haven, CT 06520-8281  
larry.samuelson@yale.edu

Michael A. Kuhn  
Department of Economics  
1285 University of Oregon  
Eugene, CA 97403  
USA  
mkuhn@uoregon.edu

A online appendix is available at <http://www.nber.org/data-appendix/w21934>

# Contents

<b>1 Introduction</b>	<b>1</b>
<b>2 Background: Starting Small</b>	<b>2</b>
2.1 Intuition . . . . .	2
2.2 Endogenous Stakes . . . . .	3
2.3 Predictions . . . . .	4
<b>3 Experimental Procedures</b>	<b>4</b>
<b>4 Results</b>	<b>5</b>
4.1 What Value of $\lambda$ Maximizes Joint Payoffs? . . . . .	5
4.2 What Value of $\lambda$ do Players Choose? . . . . .	8
4.3 Do Players Earn More When $\lambda$ is Endogenous? . . . . .	12
4.4 Is Cooperation Self-Reinforcing? . . . . .	15
<b>5 Discussion and Conclusion</b>	<b>18</b>
<b>Appendix</b>	<b>23</b>
<b>Subjects' Instructions: Random Condition</b>	<b>23</b>
<b>Subjects' Instructions: Endogenous Condition</b>	<b>29</b>

# 1 Introduction

Consider a professor who moves to a new city and must find a dentist. Knowing nothing, she selects one at random. However, she chooses to have a minor procedure done first, proceeding to major work if all goes well. She also has planned an anniversary weekend getaway with her spouse, and needs to find a babysitter. She hires a number of baby sitters for “trial runs” before choosing one for the weekend. She behaves similarly when finding a mechanic or plumber, when working with a research assistant or coauthor, and when setting the boundaries for her older daughter’s dating and use of the car.

When given the flexibility to structure relationships, it is intuitive that people start small, reserving the option to take larger risks if things go well. This paper reports an experiment examining whether players in a simple relationship—a twice-played prisoners’ dilemma—arrange the stakes across the two stages to as to start small, and whether doing so increases their payoffs.

There are many reasons why relationships might start small. The key to some of the examples in the introductory paragraph is the ability to switch partners if one’s initial experience is unsatisfactory. Section 2 notes that starting small can increase effective discount factors. This paper focuses on yet another factor, arising out of incomplete information about preferences. Kreps, Milgrom, Roberts and Wilson (1982) highlighted the role of incomplete information in the prisoners’ dilemma, examining finitely repeated prisoners’ dilemma games whose players believe their opponents may be principled types who play a cooperative strategy, such as *Tit-for-Tat*. If the interaction is sufficiently long, then it behooves others to act as if they might also be one of these principled types, leading to widespread cooperation, even if principled types are very rare (or indeed even if principled types do not exist, as long as people believe others believe others believe they *might* exist).

We believe that principled types not only exist, but are not necessarily rare. For example, experimental research points to nonnegligible proportions of people who split evenly in the dictator game or cooperate in the prisoners’ dilemma.<sup>1</sup> We see such people in our daily lives—people who make it a point to not lie, cheat, or steal. We sometimes refer to these people with backhanded complements, noting that they are “honest to a fault” or that “nice guys finish last.” However, Andreoni and Samuelson (2006) show that such principled people can play a critical role in supporting cooperation in the *twice-played* prisoners’ dilemma (in contrast to the long games of Kreps, Milgrom, Roberts and Wilson (1982)). Moreover, this effect is most pronounced in relationships that start small, but not too small. Intuitively, first-stage actions reveal information about players’ types, allowing them to more effectively exploit second-stage matches between types whose preferences support cooperation. This makes it advantageous to push stakes to the second stage. However, the ratio of stakes must not be tilted too heavily toward the second stage, lest first-stage actions become cheap talk and lose their ability to convey information. Andreoni and Samuelson (2006) conduct a twice-played prisoners’ dilemma experiment, examining treatments that differ in the (randomly set) distribution of the stakes across the two stages. They find that the optimum arrangement of stakes, in terms of joint payoffs, is to play for approximately one third of the total stakes in the first stage, reserving two-thirds for the second stage.

In this work we reproduce the Andreoni-Samuelson experimental game, but this time we allow the subjects themselves to determine the relative stakes. Will individuals choose to start small? Will they earn higher payoffs by doing so? Will they find the optimum arrangement of stakes found in the Andreoni-Samuelson experiment? Our findings are clear. Not only do the subjects gravitate toward this same allocation of stakes, which again maximizes joint payoffs, but the gains from arriving at this allocation are significantly higher when they are chosen by the players rather than controlled experimentally. That is, when one’s opponent chooses to start at about one-third of the stakes, it generates more cooperation than when those same stakes are set by the experimenter. What accounts for this difference is an open but important question. Do endogenously-chosen arrangements, such as starting small, themselves convey information, and in doing so create even greater cooperation?

---

<sup>1</sup>See Andreoni and Miller (2002) and Camerer (2003, chapter 2).

We begin our analysis of this question in the next section by providing some background on starting small and the endogenous determination of relationship stakes. We then use this to make predictions for our study. We describe our experimental procedures in section 3, present the results in section 4, and conclude in section 5.

## 2 Background: Starting Small

This line of research is not the first to examine the virtues of starting small. Schelling (1960) suggested an incremental approach to funding public goods, an idea formalized by Marx and Matthews (2000) and examined experimentally by Duffy, Ochs, and Vesterlund (2007). Watson (1999, 2002) examines infinitely-repeated prisoners' dilemma games whose stakes vary over time, identifying circumstances under which a profile of increasing stakes plays a key role in supporting cooperation. Rauch and Watson (2003) present empirical evidence that starting small plays a role in developing commercial relationships in developing countries. These papers generally are interested in developing longer-term relationships and use starting small as a means to increase the effective discount rate—as stakes rise the future becomes more important, thus disciplining impatient players. The Andreoni-Samuelson model takes a different tack, focusing on very short relationships in which discounting plays no meaningful role.<sup>2</sup>

### 2.1 Intuition

This section briefly describes the model and results of Andreoni and Samuelson (2006), counting on readers to refer to the original for details. Two players play a prisoners' dilemma, observe the outcome, and then (without discounting) play another prisoners' dilemma. Figure 1 presents the parameters of the games used here and by Andreoni and Samuelson (2006). The variables  $x_1$  and

		Stage 1		Stage 2	
		C	D	C	D
C		$3x_1$ $3x_1$	$0$ $4x_1$	$3x_2$ $3x_2$	$0$ $4x_2$
	D	$4x_1$ $0$	$x_1$ $x_1$	$4x_2$ $0$	$x_2$ $x_2$

**Figure 1: Two-Stage Prisoners' Dilemma with Variable Stake Allocation**

$x_2$  determine the stakes for which the game is played in each stage. We restrict attention to values  $0 \leq x_1 \leq 10$  and  $x_1 + x_2 = 10$ . The key variable will be the relative sizes of the stakes in the two stages, which we capture by defining  $\lambda = x_2/(x_1 + x_2)$ , so that  $\lambda$  is the fraction of total payoffs reserved for stage 2. Starting small means  $\lambda > 1/2$ .

The players in the model are heterogeneous in their willingness to be cooperative, ranging from those who never cooperate to those who always do. In particular, we suppose that each player's preferences can be characterized by a number  $\alpha$ , where an individual playing a single-shot prisoners' dilemma game will prefer to cooperate *if* they believe their opponent will cooperate with

---

<sup>2</sup>Others that develop theoretical models in which starting small optimally builds relationships include Andreoni and Samuelson (2006), Blonski and Probst (2004), Datta (1996) and Diamond (1989). Laboratory evidence on starting small is provided by Binmore, Proulx, Samuelson and Swierzbinski (1998), who investigate interactions preceded by small sunk costs, and Andreoni and Samuelson (2006). Weber (2006) uses the laboratory to confirm that coordination is more efficient in small groups that slowly build in size.

a probability at least  $\alpha$ . We say those with lower values of  $\alpha$  are “more altruistic.” The values of  $\alpha$  range from below 0 (always cooperate) to above 1 (always defect). In a single shot of the prisoners’ dilemma, there would be at least one fixed point where exactly  $\alpha^*$  fraction of the population have preference parameters less than or equal to  $\alpha^*$ , and there would exist a corresponding equilibrium in which proportion  $\alpha^*$  of the players cooperate.

We consider a twice-played prisoners’ dilemma. Call the first play stage 1 and the second play stage 2. To build intuition, think first of equal stakes across the two stages ( $\lambda = 1/2$ ). Now some people who otherwise would not cooperate in a single-shot game will cooperate in the first play of the two-stage game, in order to pool with people who have lower  $\alpha$ ’s and thereby induce their opponents to cooperate in stage 2. In equilibrium, there exists a critical point  $\alpha_1 > \alpha^*$  where all those with  $\alpha \leq \alpha_1$  will cooperate in the first stage. Moreover, observing cooperation provides good news about the opponent’s value of  $\alpha$ . This gives rise to a critical value  $\alpha_2$  such that those with  $\alpha \leq \alpha_2$  and who have experienced mutual cooperation in the first stage will also cooperate in stage 2. Importantly, in the game with equal stakes,  $\alpha_2 > \alpha^*$ .

Next consider what happens as we move stakes from the first stage to the second. This has two effects. First it increases the desire to pool with lower-type  $\alpha$ ’s in the first stage by lowering the risk of cooperating, while also increasing second-stage payoffs and hence the payoff from inducing cooperation in the second stage. We thus have a force tending to increase the incidence of mutual cooperation in the first stage and also to increase the benefits from mutual cooperation in the second stage. On the other hand, a more valuable second stage makes defecting more attractive to high  $\alpha$  types, tending to decrease cooperation in the second stage. If we assume reasonable amounts of smoothness to the distribution of  $\alpha$ ’s, then when we make a small movement away from equal stakes toward larger stakes in stage 2, the first effect will dominate—more cooperation will be seen in the first stage and the gains in payoffs in the second stage will outweigh the deleterious effects of temptation in the second stage. On net, people will be better off. As more stakes get moved to the second stage there is more pooling in the first stage, meaning that a mutually cooperative first stage is less predictive of cooperation in the second, while second-stage defecting becomes more tempting. Eventually, the marginal benefits of first stage cooperation are balanced by the marginal cost of second stage defection. Overall earnings are thus maximized by moving just the right amount of stakes from the first to the second stage.

## 2.2 Endogenous Stakes

In Andreoni and Samuelson (2006), the allocation of the stakes across the two stages is chosen randomly and is exogenous to the players, both in the theoretical model and the experimental implementation. In the experiment reported in this paper, the players choose this allocation.

The literature includes some similarly motivated studies in which the players choose some aspect of the game they are to play. The study of punishment in public goods games is a prominent example. The literature here indicates that exogenously engineered opportunities to punish can be destructive, while the endogenous adoption of delegated enforcement can be more effective.<sup>3</sup> Related work by Andreoni (2014) shows that the voluntary adoption of “satisfaction guaranteed” policies by merchants can also be useful when interactions between merchants and customers are too infrequent to build reputations. For the most part and for good reason, experimentalists work with rigidly defined games rather than allowing the structure of the interaction be determined by subjects. Peters (1995), however, develops a theory of equilibrium in markets in which multiple trading mechanisms exist, and the emergence of a dominant mechanism is endogenously determined by market participants. This work is more broadly related to the literature on

---

<sup>3</sup>On punishment see Fehr and Gächter (2000), and on its pitfalls see Nikiforakis (2008) and Rand, Dreber, Ellingsen, Fudenberg and Nowak (2009). On voluntary adoption of delegated enforcement, see Kocher, Martinsson and Visser (2012) and Andreoni and Gee (2012).

coordination on political platforms. In the political science literature, Greif and Laitin (2004) adapt the traditional theory in which institutions are defined by exogenously given parameters and endogenously determined variables by defining quasi-parameters, that is, parameters that are fixed in the short run, but variables in the long run. They apply this theory to a number of historical case studies of markets.

Endogenous group and network formation has received more attention. Recently, Charness and Yang (2014) use the laboratory to investigate how behavior, earnings and efficiency are affected by a voting procedure that allows groups participating in a public-goods game to determine their own members. They identify substantial gains in contributions from such group formation.<sup>4</sup> Ali and Miller (2013) study a theoretical model of a networked society in which the formation of each link is endogenously determined by individuals. Considered along with the study we present below, one could view this recent work as creating a framework for the development of relationships, communities and enforcement mechanisms in an environment otherwise devoid of institutions.

## 2.3 Predictions

The setting examined in this paper differs from the Andreoni-Samuelson (2006) model by allowing the players to choose the relative stakes of the two stages of the prisoners' dilemma, instead of fixing them exogenously. If the players have common priors on the distribution of preferences and are able to solve for the equilibrium, then there exists an equilibrium in which they immediately select the expected payoff maximizing allocation of stakes. Conditional on those stakes, the play should be exactly as found by Andreoni and Samuelson for the case of exogenous stakes. This will be our benchmark prediction.

Of course, experimental subjects typically do not settle on equilibria so easily. Moreover, there is no reason to believe that all subjects have equal or accurate priors on the distributions of preferences in the sample. Hence, if a player observes cooperation from an opponent in the first stage, they may attribute this choice to confusion or the need to learn equilibrium behavior, which might make players more hesitant to cooperate. We might thus expect players to achieve less cooperation than predicted by the model. On the other hand, when the stakes are chosen by the opponent, this choice itself may convey information. Choosing stakes close to the optimum may be a signal of an understanding of the underlying game, optimal strategy, and accurate beliefs, and thus may combine with a cooperative action in the first stage to send an especially strong signal of an intent to cooperate. If so, then, relative to the Andreoni-Samuelson results, one might expect that endogenously chosen stakes near the optimal amount could increase cooperation, while those farther away from the optimum could decrease payoffs. This provides our alternative prediction.

## 3 Experimental Procedures

We examine data from a total of eight experimental sessions, including five from the original Andreoni and Samuelson paper, where  $\lambda$  is chosen by the experimenter, and three new sessions where  $\lambda$  is chosen by the players themselves.<sup>5</sup> In the original Andreoni-Samuelson data, each session had 22 subjects playing 20 twice-played prisoners' dilemmas, with no player meeting the same partner twice. In the new data, two of the three sessions again had 22 subjects per session participating in 20 rounds, again with new partners, but with the subjects choosing  $\lambda$ . We will call this the *short sample*. One additional new session was extended to 40 rounds, again using 22 subjects, and this time the subjects were instructed that no two players would meet more than twice.

---

<sup>4</sup>Other notable papers on endogenous group regulation include Erhart and Keser (1999), Cinyabuguma, Page and Putterman, (2005), Page, Putterman and Unel (2005) and Ahn, Isaac and Salmon (2008,2009).

<sup>5</sup>All data was collected at the University of Wisconsin, Madison over the course of a single semester, making them comparable in terms of subject pool and timing. Copies of the experimental instructions are available in the online appendix.

We will call this the *long sample*. In the short sample we have 440 new interactions (11 pairs per round  $\times$  20 rounds per session  $\times$  2 sessions), and we also have 440 new interactions in the long sample (11 pairs per round  $\times$  40 rounds per session  $\times$  1 session).

Combining the new data with the Andreoni-Samuelson data, we can split the sample into an endogenous condition, referring to the new data in which  $\lambda$  is endogenously determined by subjects, and a random condition, referring to the original data in which the computer randomly drew  $\lambda$  from a discrete distribution ranging from zero to one with equal weight on each increment of a tenth (inclusive on both ends). The original data come from 5 sessions each involving 22 subjects playing 20 rounds of the twice repeated prisoners' dilemma, implying 1100 data points. For all side-by-side comparisons of the original and new data that follow, we exclude the new 40 round session.<sup>6</sup>

In all trials, subjects used isolated computer stations to play against a randomly matched, anonymous opponent. The prisoners' dilemma game was presented to the subjects as the "push-pull" game (Andreoni and Varian, 1999). Tokens pushed to an opponent were tripled, while tokens pulled to one's self were received at face value.

In the endogenous condition, subjects were asked explicitly for the "pull value" they wished to play for in stage 1. For example, choosing a pull value of 4 implies that in stage 1, the subjects could either pull 4 to themselves or push 12 to their partner, and in stage 2, the subjects could either pull 6 to themselves or push 18 to their partners. Therefore, a choice of 4 would correspond to a  $\lambda$  of 0.6. Both subjects were asked to submit their preferred pull value prior to each game, and the computer randomly chose one of the two submissions for use. Subjects were only told of the value of  $\lambda$  chosen, and not which player selected the value.<sup>7</sup> Subjects were paid for their performance in all games in cash following the experiment.

## 4 Results

We present our results in four parts. First (Section 4.1), we ask whether the  $\lambda$  that maximizes joint payoffs from the twice played prisoners' dilemma in the endogenous condition is similar to that in the random condition. We show that in fact they are nearly identical. Second (Section 4.2), we present evidence that subjects are indeed migrating towards the value of  $\lambda$  established as joint payoff maximizing. Third (Section 4.3), we show that behavior and payoffs, conditional on  $\lambda$ , differ under the random and endogenous conditions, with higher payoffs appearing in the endogenous conditions. Finally (Section 4.4), we explore the degree to which cooperation is self-reinforcing.

### 4.1 What Value of $\lambda$ Maximizes Joint Payoffs?

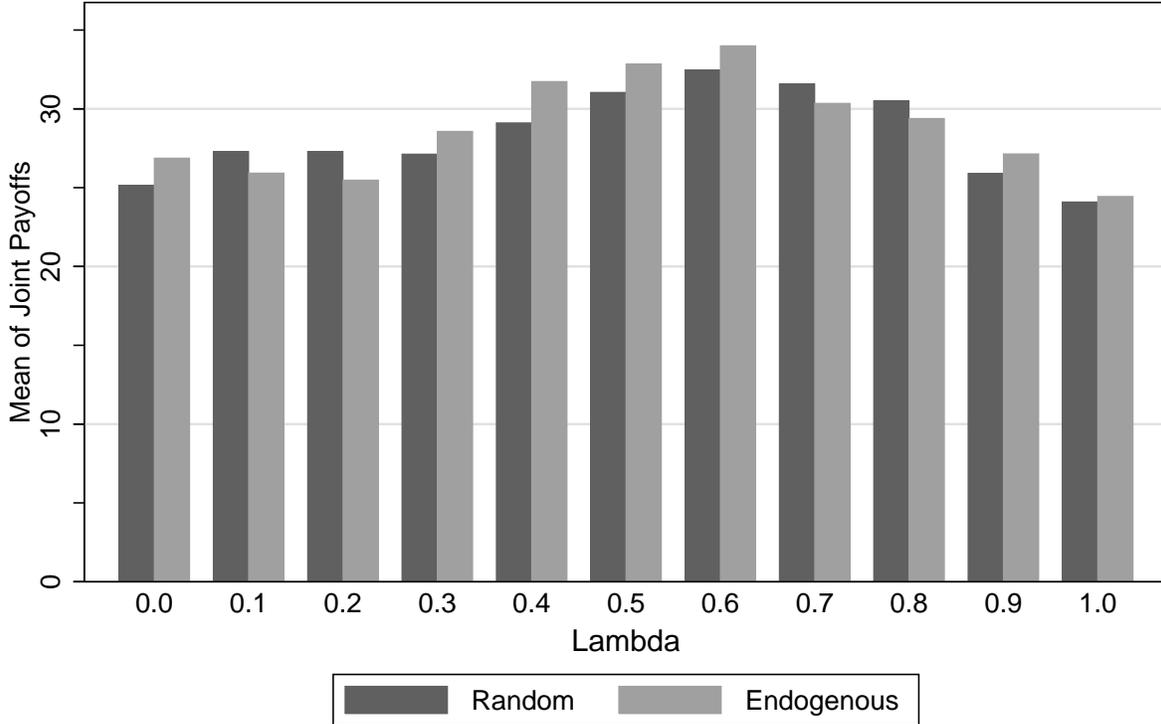
The relative stakes, measured by  $\lambda$ , are randomly chosen in the random condition and chosen by the subjects in the endogenous condition. A first inspection of the data indicates that the payoff-maximizing value of  $\lambda$  is similar in the two conditions. In particular, Figure 2 presents the mean joint payoffs from a single (two-stage) interaction separately by  $\lambda$  and by condition. A value of  $\lambda = 0.6$  appears to be optimal for both random and endogenous conditions.

To examine this more precisely, we follow Andreoni and Samuelson (2006) and estimate joint payoffs,  $\pi$ , as a cubic polynomial of  $\lambda$ , conditional on a round fixed effect,  $\gamma_t$ . We then find the value of  $\lambda$  that maximizes this polynomial. We call this first

---

<sup>6</sup>Results from this long session are very similar to those from the short sessions, so we only present results from the long sessions when we wish to focus on issues specific to the experiment length.

<sup>7</sup>For instance, if both players chose the same value, this fact was never revealed. This part of the design was intended to keep the degree of information about one's partner as similar as possible across all plays of the game. With this design, no partner whose chosen  $\lambda$  is used will know the value of  $\lambda$  chosen by the other player. This fact is a constant across all games.



**Figure 2: Payoffs by Value of  $\lambda$  and Condition**

specification CP, for cubic polynomial:

$$\text{CP : } \pi_{i,t} + \pi_{j,t} = \gamma_t + \beta_1 \lambda_{k(i,j),t} + \beta_2 \lambda_{k(i,j),t}^2 + \beta_3 \lambda_{k(i,j),t}^3 + \epsilon_{i,j,t},$$

where  $i$  and  $j$  denote two individuals paired in round  $t$  and  $k(i, j)$  is an index indicating whether individual  $i$ 's or  $j$ 's value of  $\lambda$  is chosen.

In principle, the choice of  $\lambda$  may not depend on individual preferences for cooperation because, regardless of one's plans for subsequent game play, it is desirable to induce one's partner to cooperate. In reality, individual characteristics are likely to play a role in determining the chosen  $\lambda$ , and these same personal characteristics are likely to influence how people play the game once  $\lambda$  is determined and thus how much they earn from playing. To account for this, we take an individual fixed effect approach in specification FE:

$$\text{FE : } \pi_{i,t} + \pi_{j,t} = \theta_i + \theta_j + \gamma_t + \beta_1 \lambda_{k(i,j),t} + \beta_2 \lambda_{k(i,j),t}^2 + \beta_3 \lambda_{k(i,j),t}^3 + \epsilon_{i,j,t},$$

where  $\theta_i$  and  $\theta_j$  are individual-specific constants for both players in a pairing. This approach should improve our estimate of the relationship between  $\lambda$  and payoffs in the endogenous condition. Given both the individual and round constants, any remaining confounding endogeneity of  $\lambda$  must be within-individual, time-varying covariance between the choice of  $\lambda$  and the cooperation decision.

Table 1 shows the CP and FE specifications side-by-side, applied to only the endogenous condition. We test for the joint equality of all shared coefficients across the two models (including round fixed effects). While we reject the null hypothesis that

the model coefficients are jointly equal across specifications, the non-linear combination of the  $\lambda$  coefficients yields very similar estimates  $\lambda^*$  of the payoff-maximizing  $\lambda$ , and we do not reject equality of the payoff maximizing  $\lambda$  across specifications.

**Table 1: Relationship between  $\lambda$  and Payoffs in the Endogenous Condition**

	Model	
	CP	FE
$\lambda$	3.802 (13.289)	1.487 (12.166)
$\lambda^2$	51.013 (33.097)	40.451 (29.908)
$\lambda^3$	-57.439** (22.906)	-42.722** (20.550)
$H_0$ : CP Terms Jointly = FE Terms	$\chi^2(22)=37.35^{**}$ ( $p = 0.02$ )	
Payoff-maximizing $\lambda$ : $\lambda^*$	0.627 (0.025)	0.649 (0.030)
$H_0$ : $\lambda_{CP}^* = \lambda_{FE}^*$	$\chi^2(1) = 0.97$ ( $p = 0.33$ )	
$N$	440	440

\* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . Standard errors are reported in parentheses. The payoff-maximizing  $\lambda$  is a non-linear combination of the three coefficient estimates obtained using the quadratic formula on the derivative of the implied cubic polynomial. The comparison between the CP and FE coefficients is an application of Hausman's (1978) over-identification test. It tests for the equality of all parameters common to both models.

Based on the results of the test of  $\lambda$  coefficients across models, we proceed by using the more conservative specification FE to test for interactions between condition and the effect of  $\lambda$ . Results are found in Table 2. Joint payoffs are maximized within the same range in both conditions. Restricting attention to the last 10 rounds, we see a slight drop in the payoff-maximizing value  $\lambda^*$  from 0.68 in the random condition to 0.62 in the endogenous condition. However these estimates are not significantly different at typical confidence levels. Also notice that the maximum is more tightly defined in the second half of the experiment, but only in the endogenous condition; the standard error of the maximum estimate is about 3.5 times larger in the first ten rounds than in the last ten.

An important consideration in the regressions presented above is that the unit of analysis is the game. Given that players are switching partners each game, it is impossible to apply any standard form of clustering of the standard errors. The implication of this is that the standard errors are likely to be biased downward.<sup>8</sup> If instead we took individual payoffs, rather than joint payoffs, as our unit of analysis, then we would be double counting by including a player's actions in the payoffs of two people.<sup>9</sup> In this section, the issue is unimportant since we do not find significant across-condition differences. However, in section 4.3, where we do identify such differences, we employ two alternatives to the unadjusted standard errors.

<sup>8</sup>By assuming that each pairing represents a statistically independent observation, we are over-counting the data to some degree.

<sup>9</sup>Doing so, and clustering at the individual level, would impose the odd restriction that the unobserved error term (which is by necessity the same in each count of the same observed game) is correlated with both players' other error terms, but that each player's set of error terms are uncorrelated.

**Table 2: Relationship between  $\lambda$  and Payoffs across Conditions, FE Model**

	Sample Restriction		
	All Rounds	Rounds 1-10	Rounds 11-20
$\lambda$	-1.408 (6.474)	-12.623 (10.785)	-3.205 (6.979)
$\lambda * 1(\textit{endogenous} = 1)$	4.186 (13.453)	19.080 (22.073)	5.619 (15.293)
$\lambda^2$	43.823*** (15.414)	66.310*** (25.959)	49.563*** (16.487)
$\lambda^2 * 1(\textit{endogenous} = 1)$	-6.602 (32.776)	-57.092 (55.101)	2.515 (36.201)
$\lambda^3$	-44.768*** (10.041)	-59.212*** (17.059)	-46.765*** (10.660)
$\lambda^3 * 1(\textit{endogenous} = 1)$	4.819 (22.227)	46.462 (37.648)	-10.990 (24.269)
$H_0$ : Endogenous Terms = 0	$F(3, 1361) = 0.46$ ( $p = 0.71$ )	$F(3, 601) = 1.38$ ( $p = 0.25$ )	$F(3, 601) = 1.53$ ( $p = 0.20$ )
Random Payoff-max. $\lambda_r^*$	0.636 (0.020)	0.635 (0.030)	0.673 (0.018)
Endogenous Payoff-max. $\lambda_o^*$	0.656 (0.033)	0.717 (0.145)	0.623 (0.027)
$H_0 : \lambda_r^* = \lambda_o^*$	$\chi^2(1) = 0.28$ ( $p = 0.60$ )	$\chi^2(1) = 0.31$ ( $p = 0.58$ )	$\chi^2(1) = 2.28$ ( $p = 0.13$ )
$R^2$	0.534	0.575	0.599
$N$	1540	770	770

Notes: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . Standard errors are in parentheses under the estimates. The payoff-maximizing  $\lambda$  is a non-linear combination of the three coefficient estimates obtained using the quadratic formula on the derivative of the implied cubic polynomial.

## 4.2 What Value of $\lambda$ do Players Choose?

Here we first ask whether subjects see and learn the strategy of starting small. We then look more specifically at the  $\lambda^*$  found in Section 4.1 and ask whether subjects in the endogenous condition come to choose this value with greater frequency over the course of the study.

We first partition our sample into three intervals. We call rounds 1 to 6 the beginning, rounds 7-14 the middle, and rounds 15-20 the end. We further sort subjects based on their choice of  $\lambda$ . Any subject whose average choice is less than 0.5 is said to start large, while if it is greater than 0.5 they are said to start small. If the average exactly equals 0.5 we say they start even. Table 3, Panel A presents the proportions of individuals who choose to start small, even, or large. Interestingly, in the beginning a majority starts large, and by the end the pattern has flipped with a majority starting small. We then look separately at those who started small in the beginning, and those who started large in the beginning. Both groups gravitate to starting small by the end, and those who

started started small at the beginning do so to even a greater degree.<sup>10</sup>

**Table 3: Evolution of Sub-Sample Sizes over Time**

Group, Rounds	Start Small	Even	Start Large
Panel A: Short Sample			
Unconditional Groups			
Beginning (Rounds 1-6)	0.36	0.02	0.61
Middle (Rounds 7-14)	0.48	0.07	0.45
End (Rounds 15-20)	0.54	0.09	0.36
Conditional on Starting Small in Beginning			
End (Rounds 15-20   Rounds 1-6 = SS)	0.63	0.19	0.19
Conditional on Starting Large in Beginning			
End (Rounds 15-20   Rounds 1-6 = SL)	0.52	0.04	0.44
$N = 44$ per round			
Panel B: Long Sample			
Unconditional Groups			
Beginning (Rounds 1-6)	0.45	0.14	0.41
Middle (Rounds 17-24)	0.64	0.09	0.27
End (Rounds 35-40)	0.64	0.09	0.27
Conditional on Starting Small in Beginning			
End (Rounds 35-40   Rounds 1-6 = SS)	0.70	0.00	0.30
Conditional on Starting Large in Beginning			
End (Rounds 35-40   Rounds 1-6 = SL)	0.78	0.00	0.22
$N = 22$ per round			

Do our subjects in the long sample continue to learn to start small after round 20? To do the same analysis in the 40-round sample, we use intervals of the same number of rounds as in the 20-round analysis to maintain comparability in classifying choices. The beginning runs from rounds 1-6, the middle from 17-24 and the end from 35-40. The results of the analysis are presented in Table 3, Panel B and corroborate what we observed in the shorter sample. Whereas 45% of the sample starts small in the beginning, 64% of the sample started small in the end, and although 41% of the sample started large in the beginning, only 27% started large in the end. While the the magnitude of the shift towards starting small is larger in the long sample, it is worth nothing that starting large is less prominent overall, even at the very beginning of play.<sup>11</sup>

<sup>10</sup>Paired  $t$ -tests of the short sample frequencies in Table 3, Panel A reveal a marginally significant difference between starting small and large in the beginning rounds ( $p = 0.09$ ) and a more robust difference between starting small and large in the end rounds conditional on starting small in the beginning rounds ( $p = 0.05$ ). All other comparisons are not significant at conventional levels.

<sup>11</sup>Paired  $t$ -tests of the long sample frequencies in Table 3, Panel B indicate that the fraction starting small is significantly greater than the fraction starting large in the middle and end with  $p = 0.07$  in both cases. The comparisons of starting large and small in the end rounds conditional on behavior in the beginning rounds are limited by very small sample size, but nonetheless the comparison conditional on starting large in the beginning is on the margin of statistical significance ( $p = 0.10$ ).

Next we present a variety of estimates of the trends in the choice of  $\lambda$  throughout the sessions. We start with  $\lambda$  itself, to ascertain an overall directional trend. We also do this in a first-differenced specification. Next, we consider the absolute deviation of  $\lambda$  from 0.656, our estimate of the payoff-maximizing  $\lambda$  over the course the short sample, to see if individuals are getting closer to that value over time. In these first three specifications, observations of  $\lambda = 0$  and  $\lambda = 1$  are excluded because the regression uses the cardinal information in  $\lambda$ .<sup>12</sup> The last two specifications do not require this exclusion. As an analogue to Table 3, we use an indicator variable for starting small as an outcome variable. Lastly, we collapse the data to the round level and calculate a round-specific Herfindahl index ( $H$ -index) that measures how “monopolistic” the market for  $\lambda$  values is. This approach is designed to assess whether learning and convergence happen over the course of a session. The results are in Table 4, with standard errors clustered by individual in all specifications except for the  $H$ -index specification.

**Table 4: Time Trends in  $\lambda$  Choice**

	Measure of $\lambda$				
	$\lambda_{i,t}$	$\lambda_{i,t} - \lambda_{i,t-1}$	$ \lambda_{i,t} - 0.656 $	$Pr(\lambda_{i,t} \in [0.6, 0.9])$	$H$ -index
Panel A: Short Sample					
Round	0.005*** (0.002)	0.002** (0.001)	-0.003** (0.001)	0.007 (0.004)	0.002** (0.001)
Constant	0.456 (0.029)	-0.018 (0.010)	0.265 (0.018)	0.306 (0.053)	0.107 (0.007)
$N$	787	705	787	880	20
Panel B: Long Sample					
Round	0.003** (0.001)	0.001*** (0.000)	-0.001* (0.000)	0.007** (0.003)	-0.000 (0.000)
Constant	0.499 (0.027)	-0.016 (0.007)	0.214 (0.010)	0.350 (0.061)	0.149 (0.010)
$N$	804	743	804	880	40

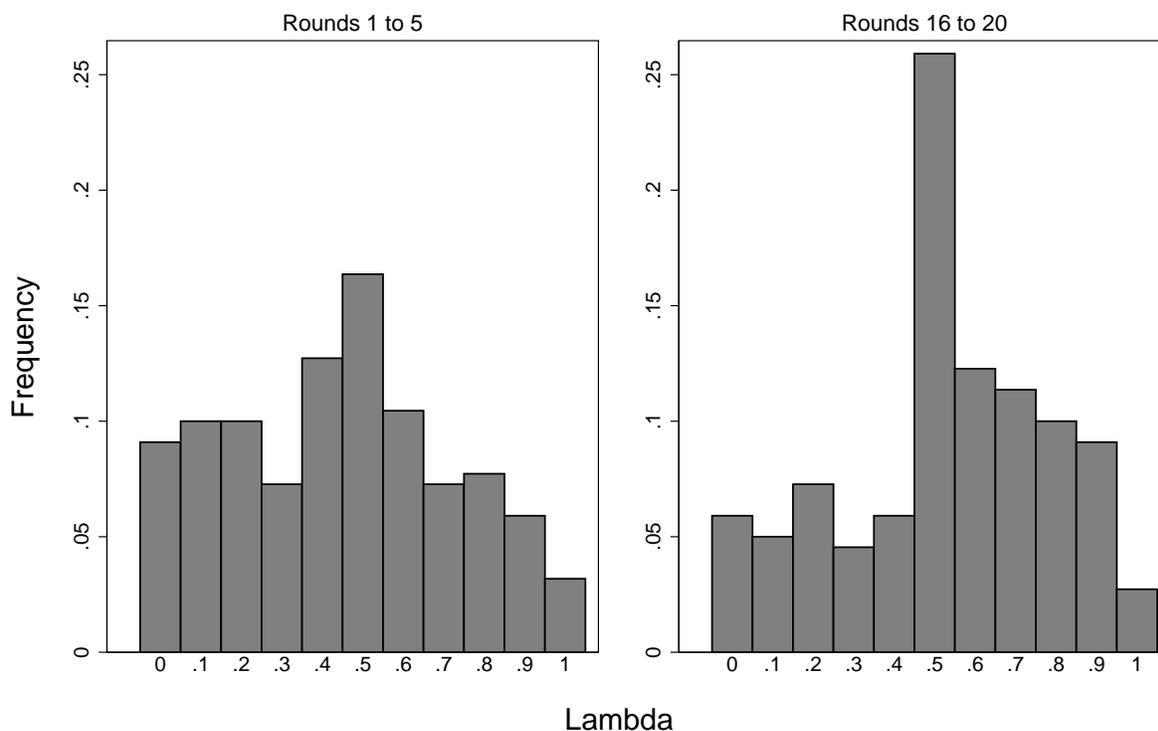
Notes: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . Standard errors are in parentheses under the estimates and clustered by individual in all specification except for the  $H$ -index, in which case observations are at the round level. In calculating the  $H$ -index, we pool  $\lambda = 0$  and  $\lambda = 1$  in order to avoid over-estimating the degree of choice dispersion.

We find, in both the short and long samples, that the mean choice of  $\lambda$  is increasing slowly and significantly over time. In the short sample, the predicted  $\lambda$  rises from 0.456 to 0.556 from round 1 to round 20. In the long sample, the predicted value rises from 0.499 in round 1 to 0.559 in round 20 to 0.619 in round 40. In both samples,  $\lambda$  is initially falling from round to round, as indicated by the negative constant in the first-difference specifications, but the positive and significant coefficient on the round variable indicates that the negative trend attenuates and then reverses to a growth trend. In the short sample, this crossover is predicted to occur after 9 rounds and after 16 rounds in the long sample. Again in both samples, the gap between choices and the payoff-maximizing choice is shrinking significantly over time. The magnitude of the round coefficient in the long sample is smaller, reflecting the fact that choices in the later part of that session are relatively stable. Consistent with the evidence from Table 3, the share of individuals starting small is increasing over both short and long samples, although the round coefficient in the small sample is not significant at conventional levels ( $p = 0.11$ ). Our measure of general learning, the  $H$ -index, grows significantly over time in the short sample.

<sup>12</sup>Not only are these extreme values cardinally ambiguous with respect to the other  $\lambda$  values, but with respect to each other as well.

We do not find this result in the long sample.<sup>13</sup>

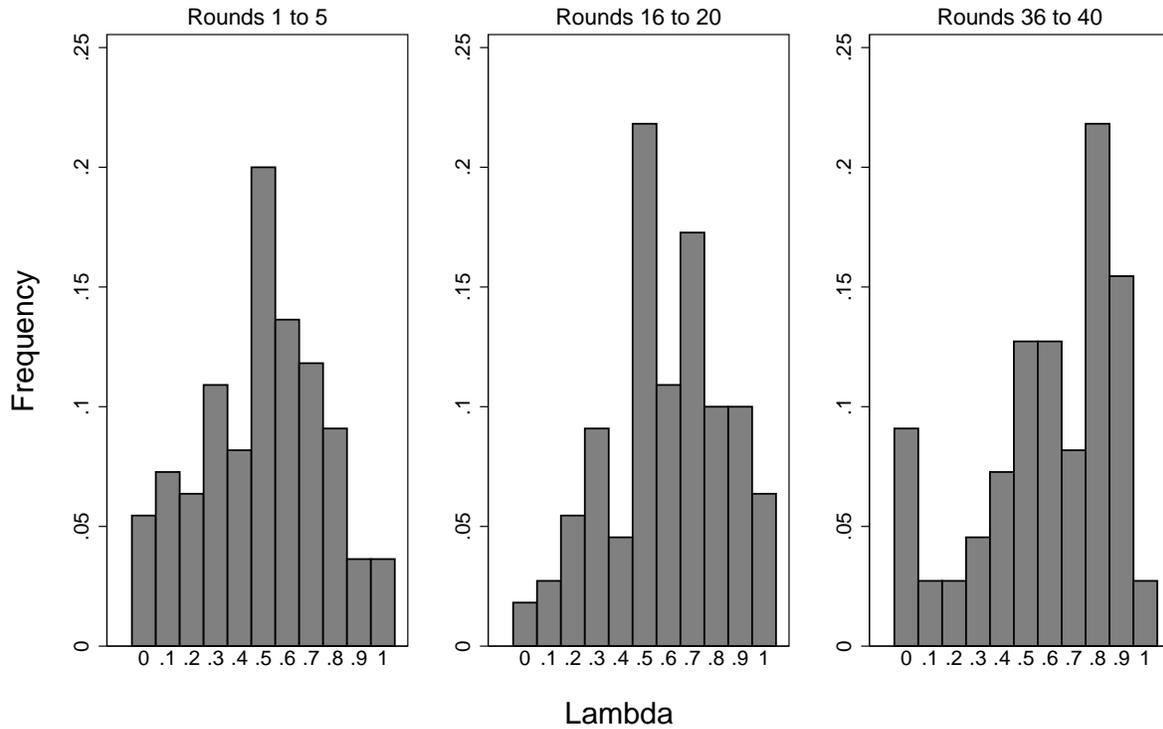
Our final approach is to focus on the changes to the full distribution of choices over the course of the short and long samples. Figure 3 shows histograms of  $\lambda$  choices for the first and last 5 rounds in the short sample, while Figure 4 shows the first 5, last 5 and rounds 16-20 (corresponding to the last 5 in the short sample) in the long sample. Both figures indicate a shift of mass from the left side of the distribution (start large) to the right side of the distribution (start small) over time. Whereas starting even is still relatively common by the end of the short sample, other values in the start small region overtake it in frequency by the end of the long sample. We use Kolmogorov-Smirnov tests to assess whether the distributions in Figures 3 and 4 differ from one another.<sup>14</sup> In the short sample, the distribution of  $\lambda$  choices in rounds 16-20 is significantly to the right of the distribution of choices in rounds 1-5 ( $D = 0.20, p < 0.01$ ), indicating movement towards starting small. In the long sample, we find a shift towards starting small in rounds 16 to 20 and rounds 36 to 40 relative to rounds 1-5 ( $D = 0.15, p = 0.07$  and  $D = 0.24, p < 0.01$  respectively). The continued shift towards starting small from rounds 16-20 to 36-40 is not significant at conventional levels ( $D = 0.14, p = 0.13$ ). Testing across short and long samples, we do not reject the equality of distributions in rounds 1-5 ( $D = 0.11, p = 0.25$ ). However, starting small is significantly more frequent in the long sample by rounds 16-20 than in the same rounds in the short sample ( $D = 0.22, p < 0.01$ ), and thus we also find significantly more starting small in rounds 36-40 of the long sample than in rounds 16-20 of the short sample ( $D = 0.26, P < 0.01$ ).



**Figure 3:  $\lambda$  Choices by Round in the Short Sample**

<sup>13</sup>Given our 10 distinct values of  $\lambda$  (we treat 0 and 1 identically), the minimum value for the index is 0.10. Part of the reason the H-index was not a stronger measure is the fact that  $\lambda$ s of 0 and 1 were disproportionately chosen, especially early in the study, perhaps because they are more focal. With repetition, these extremes became less concentrated as the intermediate points became more concentrated, which understated the change in the desired direction.

<sup>14</sup>Aggregated up to 5-round bins, testing the distributions against the hypothesis of uniformity rejects the null in all circumstances.



**Figure 4:  $\lambda$  Choices by Round in the Long Sample**

### 4.3 Do Players Earn More When $\lambda$ is Endogenous?

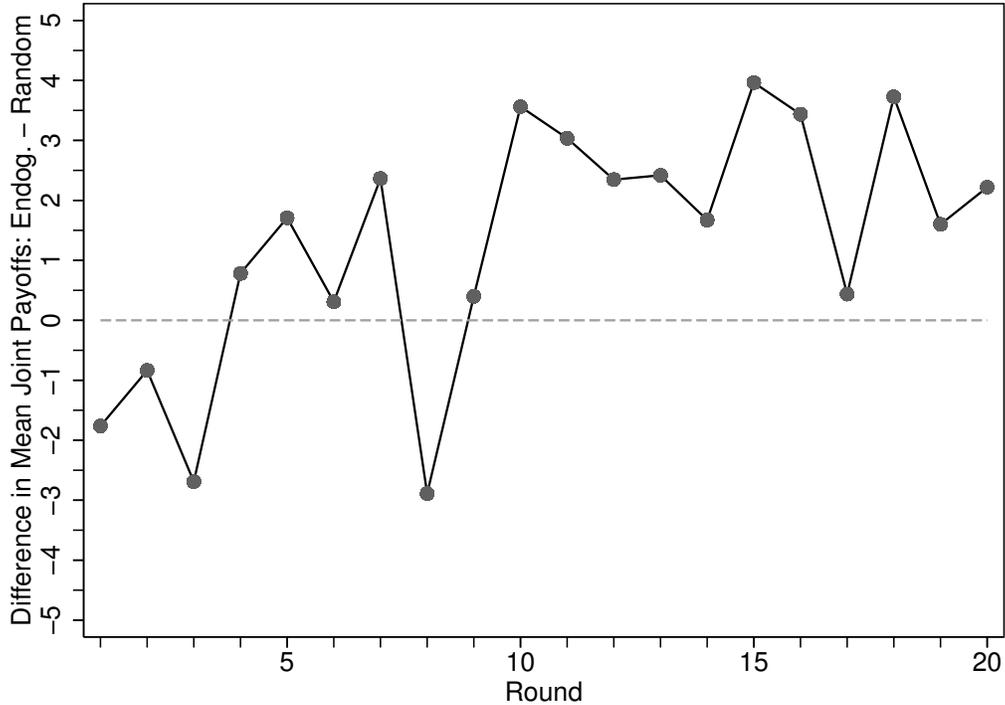
Having established that the optimal choice of  $\lambda$  is just above 0.6 and that subjects migrate in the direction of that choice over the course of the experiment, we should be able to establish that subjects earn more in the endogenous condition than in the random condition, in the later rounds of the experiment.<sup>15</sup>

Figure 5 presents the difference in mean joint payoffs across conditions for each round. The figure clearly shows a separation between endogenous and random conditions that emerges over time. While we have one test statistic for this difference, for completeness we present it with three versions of the  $p$ -values for each test. The first,  $p_1$ , treats all games within and across rounds as independent observations, the second,  $p_2$ , uses standard errors clustered at the individual level (which necessitates counting each game twice), and the third,  $p_3$ , uses a conservative manual scaling of the standard error by a factor that changes the implied sample size from the total number of games to the total number of subjects.<sup>16</sup>

Regressing joint payoffs on a condition dummy demonstrates that average payoffs are higher overall in the endogenous condition by 1.29 tokens per game. Results in columns (1) and (4) of Table 5 are for all rounds and the last 10 rounds, respectively. This represents a 4.5% gain over the random condition, and is significant according to  $p_1$  and  $p_2$  measures, but not  $p_3$ . This difference is driven by differences in the second half of the sessions. During the first 10 rounds, the mean payoff difference across conditions is an insignificant 0.09 tokens, but during the last 10 rounds, the mean payoff difference is a strongly significant 2.49

<sup>15</sup>Given monotonicity about the optimum.

<sup>16</sup>Since the standard error of the mean is the sample standard deviation of the mean divided by the square root of the number of observations, we multiply it by the square root of the number of games and then divide it by the square root of the number of subjects as an adjustment. For example, when using data from all rounds in both conditions, there are 1540 games played by 154 subjects. Thus, the adjustment in this case implies multiplying our standard errors by  $\sqrt{10}$ .



**Figure 5: Excess Payoffs in Endogenous Condition over Random Condition by Round**

tokens. This represents a 10.2% gain over the random condition, significant according to all three calculations of the standard errors.

These excess returns lead us to examine whether there are payoff differences *conditional* on  $\lambda$ . In other words, once  $\lambda$  is chosen, does it matter if that choice came from a computer or a player? We regress joint payoffs on a condition dummy while controlling for  $\lambda$ . Results are in columns (2) and (5) of Table 5. The estimate of the all-round mean payoff difference in this case is 1.32 tokens higher in the endogenous condition, which is significantly different from zero according to  $p_1$  and  $p_2$ . As with the unconditional condition effect, this conditional condition effect is only a feature of the final ten rounds: an insignificant 0.16 tokens in the first ten and a strongly significant 2.46 tokens in the last ten. However, we've already demonstrated that  $\lambda$  has a non-linear effect on payoffs, which is not captured in columns (2) and (5). To deal with this, we adopt a cubic polynomial as the control for  $\lambda$ , with results in columns (3) and (6) of Table 5. In this specification, there are only significant differences across conditions in the final ten rounds. The difference of 1.32 tokens in the latter half of the study is significant using  $p_1$  and  $p_2$ .

Next, we allow for the endogenous effect to vary with  $\lambda$ . We speculated earlier that on the one hand,  $\lambda$  values far outside the optimal range could be warning signs, and we might expect to see a negative condition effect, and on the other that values near the optimum could induce more cooperation. We find support for this hypothesis. Figure 6 reproduces Figure 2, but reporting for only the last ten rounds. At central values of  $\lambda$  of 0.4, 0.5, and 0.6 we find significantly higher payoffs in the endogenous condition, while at non-central values of  $\lambda$  equal to 0.1, 0.2, and 0.8 we find significantly lower payoff in the endogenous condition.<sup>17</sup>

<sup>17</sup>The differences at  $\lambda$  of 0.4 and 0.5 are significant using all three  $p$ -value approaches, while 0.6 is significant for  $p_1$  and  $p_2$ . For  $\lambda = 0.1, 0.8$  we find significance only using  $p_2$  and for  $\lambda = 0.2$  we find significance using all three approaches. It is important to note here, however, that  $p_3$  is not a conservative bound on the standard error in this case because often we have more individuals participating in games with a certain  $\lambda$  value than games with that value overall. Therefore,  $p_1$  and  $p_2$  are actually more conservative than  $p_3$  in this case.

**Table 5: Payoff Differences across Conditions**

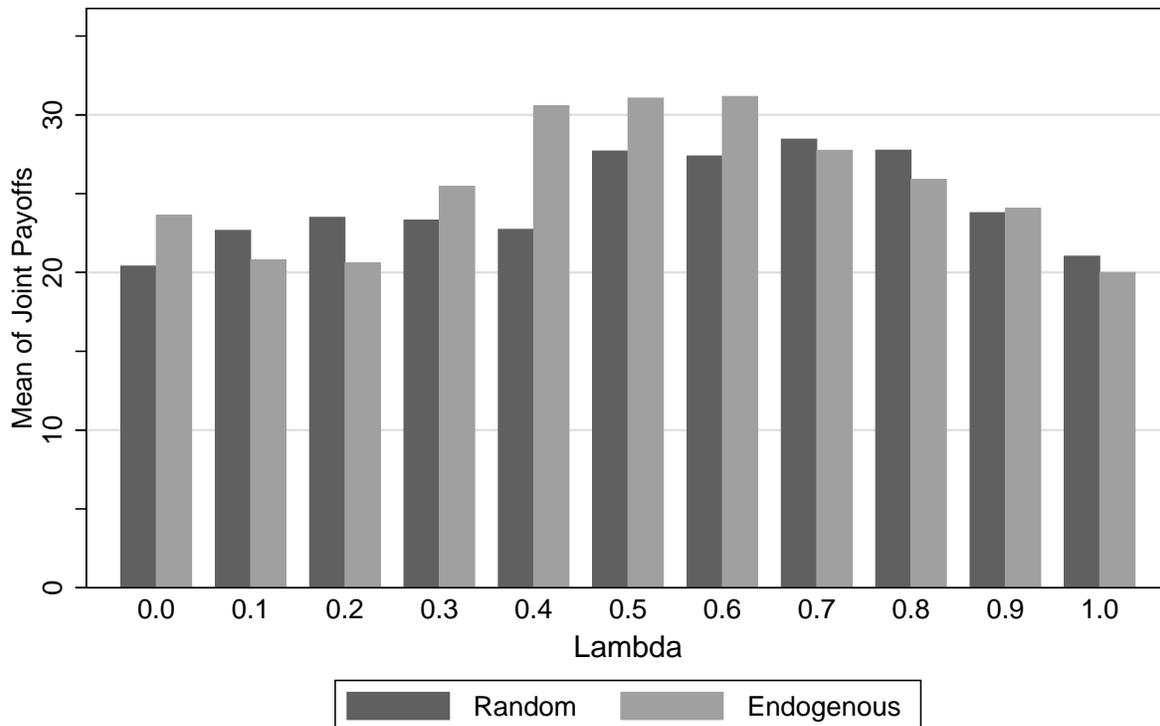
	Joint Payoff (tokens)					
	All Rounds			Rounds 11-20		
	(1)	(2)	(3)	(4)	(5)	(6)
1( <i>endogenous</i> = 1)	1.290 <sup>††</sup>	1.318 <sup>††</sup>	0.468	2.485 <sup>†††</sup>	2.458 <sup>†††</sup>	1.318 <sup>††</sup>
	(0.579)	(0.578)	(0.567)	(0.569)	(0.565)	(0.539)
	$p_1 = 0.03$	$p_1 = 0.02$	$p_1 = 0.41$	$p_1 < 0.01$	$p_1 < 0.01$	$p_1 = 0.02$
	$p_2 = 0.10$	$p_2 = 0.09$	$p_2 = 0.53$	$p_2 < 0.01$	$p_2 < 0.01$	$p_2 = 0.06$
	$p_3 = 0.48$	$p_3 = 0.47$	$p_3 = 0.79$	$p_3 = 0.05$	$p_3 = 0.05$	$p_3 = 0.27$
Constant	38.374	27.596	25.573	24.442	22.950	20.865
	(0.309)	(0.528)	(0.773)	(0.304)	(0.522)	(0.754)
Control for						
$\lambda?$	No	Yes	Yes	No	Yes	Yes
$\lambda^2, \lambda^3?$	No	No	Yes	No	No	Yes
$N$	1540	1540	1540	770	770	770

Notes: <sup>†</sup> $p_1 < 0.10$ , <sup>††</sup> $p_1, p_2 < 0.10$ , <sup>†††</sup> $p_1, p_2, p_3 < 0.10$ . Unadjusted standard errors are in parentheses under the estimates, followed by the three different  $p$ -value specifications: unadjusted, cluster-adjusted with double counting and manually scaled. Moving across columns from (1) and (4) to (2) and (5) adds a linear control for lambda and moving to columns (3) and (6) adds a cubic polynomial control for lambda.

The most direct test of whether the choice of  $\lambda$  is interpreted as a signal is to leverage the fact that when a subject's choice of  $\lambda$  is implemented, its revelation is largely uninformative. For example, if a subject chooses  $\lambda = 0.6$  and then observes  $\lambda = 0.6$  implemented, their interpretation should be that in all likelihood their choice was implemented, albeit with a small probability that their partner chose the same. However, when a subject chooses  $\lambda = 0.6$  and then observes  $\lambda = 0.5$  implemented, they know exactly their partner's choice and that their strategic intentions are similar. The same is true for extreme values of  $\lambda$ . Our signaling hypothesis is that subjects who chose central values of  $\lambda$  (0.4, 0.5, 0.6) in games with a central value of  $\lambda$  implemented should be more likely to cooperate when their choice is not selected because the revelation of  $\lambda$  is informative. On the other hand subjects who chose extreme values of  $\lambda$  (all non-central values) in games with an extreme value of  $\lambda$  implemented should be less likely to cooperate when their choice is not selected for the same reason. We regress an indicator for an individual cooperating in the first stage on four mutually exclusive indicators:

- an individual's central choice of  $\lambda$  was implemented: *central\_used*
- an individual's central choice of  $\lambda$  was not implemented and a different central  $\lambda$  was implemented: *central\_unused*
- an individual's extreme choice of  $\lambda$  was implemented: *extreme\_used*
- an individual's extreme choice of  $\lambda$  was not implemented and a different extreme  $\lambda$  was implemented *extreme\_unused*

We expect to find that the coefficient on *central\_unused* is greater than that on *central\_used* and the coefficient on *extreme\_unused* is less than that on *extreme\_used*. Relative to the excluded group (subjects aware of their partner's  $\lambda$  choice being misaligned with theirs), we expect that the sign of the *central\_unused* coefficient to be positive and the coefficient on *extreme\_unused* to be negative. We use a cubic polynomial of  $\lambda$ , individual fixed effects and game dummy variables as controls. Results are presented in Table 6.



**Figure 6: Payoffs and Lambda across Conditions, Last 10 Rounds Only**

Using all rounds in the endogenous condition, our central  $\lambda$  indicator variables have no relationship with first stage cooperation and thus there is no difference in behavior across groups. However, there is a significant negative relationship between *extreme\_unused* and cooperation which leads to a large and significant difference; in games with extreme  $\lambda$  values in which both players choose such values, cooperation is 12% less likely for a player that is aware their partner chose an extreme value as well. Looking only at the last 10 rounds, we see a large positive relationship between *central\_unused* and cooperation. While the difference by group is large—in games with central  $\lambda$  values in which both players choose such values, cooperation is 18% more likely for a player that is aware their partner chose a central value as well—it is not significant at conventional levels ( $p = 0.16$ ). In extreme value games, the findings are similar in the last 10 rounds; observing a partner’s  $\lambda$  is associated with a decrease in the likelihood of cooperation by 13%.

#### 4.4 Is Cooperation Self-Reinforcing?

Andreoni and Samuelson (2006) explain cooperation in the twice repeated prisoners’ dilemma as rationally emerging from a model of innate preferences for cooperation. Could individuals learn about their own preferences for cooperation through their experiences participating in socially beneficial actions? In other words, does one learn about their warm-glow enjoyment of cooperation by “accidentally” having a successful cooperative experience? Our data offer a unique ability to ask whether participating in a successful cooperation in the past reinforces cooperative behavior (i.e., is “habit forming”) using exogeneity in the determination of  $\lambda$ . We estimate the causal impact of having cooperated in the previous round on the likelihood of cooperating in the present round. Furthermore, we determine whether this reinforcement effect is stronger in the random condition or the endogenous condition.

Isolating the causal impact of cooperation in the past on cooperation in the future requires finding random variation in whether

**Table 6: Effect of  $\lambda$  in the Endogenous Condition on Cooperation**

	First Stage Cooperation	
	All Rounds (1)	Rounds 11-20 (2)
<i>central_used</i>	-0.017 (0.063)	0.047 (0.084)
<i>central_unused</i>	-0.053 (0.075)	0.230** (0.110)
Difference	0.036 (0.073)	-0.183 (0.129)
<i>extreme_used</i>	0.028 (0.043)	0.054 (0.064)
<i>extreme_unused</i>	-0.091* (0.048)	-0.074 (0.077)
Difference	0.119*** (0.042)	0.128* (0.069)
<i>N</i>	880	440

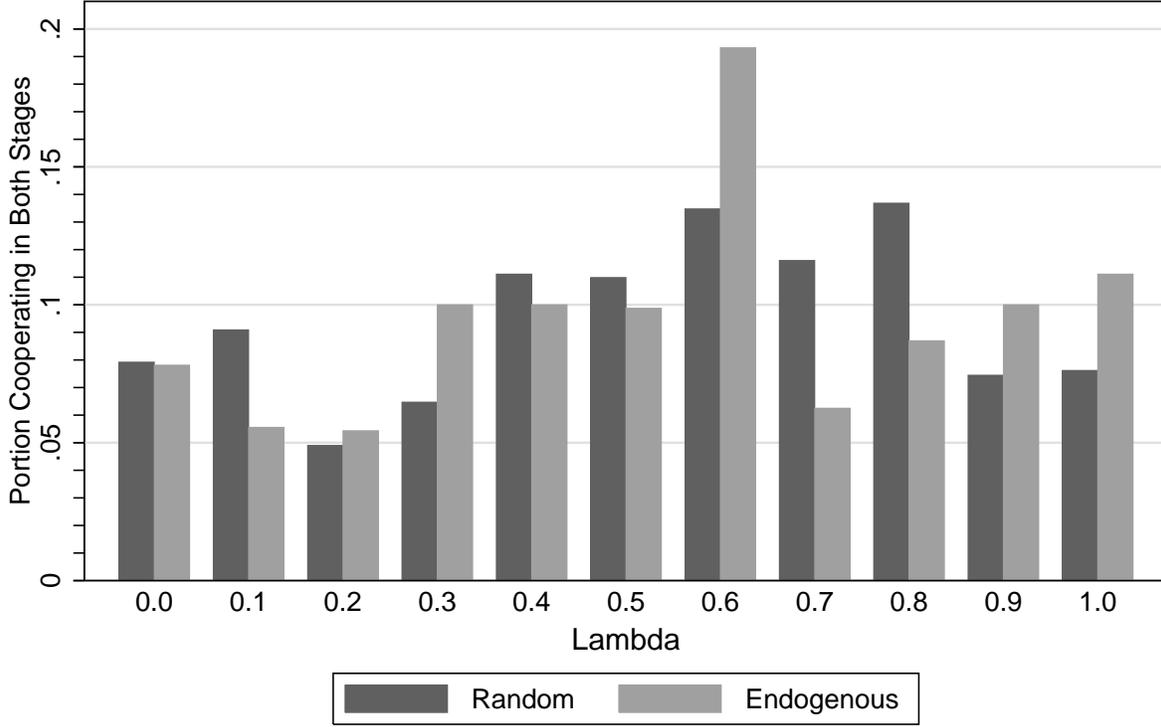
Notes: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . Standard errors, clustered by individual are in parentheses under the estimates. Both specifications include a cubic polynomial of  $\lambda$ , individual fixed effects and game dummy variables as controls

an individual chose to cooperate in the past. A nice example comes from Fujiwara, Meng, and Vogl (2013), in which weather events alter the transactions costs of voting. Using instrumental variables, this allows the researchers to identify the causal impact of voting in the past on voting in the future. In the case of our study, we need an instrumental variable for cooperation in any given round that will serve the role of the weather shocks to voting costs: what random source of variation affects the decision to cooperate? We use  $\lambda$  for this. Here  $\lambda$  can be thought of as a cost of cooperation, and random variation in  $\lambda$  can thus lead to random variation in cooperation.

The difference in how  $\lambda$  is determined between the random and endogenous conditions requires that we use two different approaches to using it as an instrument for cooperation. In the random condition, we use  $\lambda$  as an instrument for whether an individual cooperates in both stages of a round, controlling for the first stage behavior of their partner in that round. In the endogenous condition, we limit the sample to subjects in rounds that encounter a  $\lambda$  that they did not choose. We then use  $\lambda$  as an instrument in the same way. When we run the second stage of the instrumental variable regressions—cooperation in both stages of the current game regressed on cooperation in both stages of the prior game, adjusting for the endogeneity of cooperation in the prior game—we add the additional control of  $\lambda$  in the current game and their partner's behavior in the first stage of the current game. In the endogenous condition, we also control for whether an individual's choice of  $\lambda$  was implemented in the current game.

Maximizing the relevance of our instruments requires a different functional form across conditions. Figure 7 shows the relationship between  $\lambda$  and the likelihood of cooperating in both stages of a round. In the random condition, the cubic approximation used earlier to estimate the relationship between  $\lambda$  and joint payoffs fits well. In the endogenous condition, an indicator variable for whether  $\lambda$  is selected to be its nearest-to-cooperation-optimal value of 0.6 appears to be a better predictor of cooperation due to

the large spike in likelihood there and the noisy relationship elsewhere.



**Figure 7: Cooperation and  $\lambda$  across Conditions**

Our first-stage IV specifications are

$$\text{IV-R1: } 1(C_{i,t-1}^1, C_{i,t-1}^2 = 1) = \alpha_i + \gamma_{t-1} + \beta_1 \lambda_{t-1} + \beta_2 \lambda_{t-1}^2 + \beta_3 \lambda_{t-1}^3 + \delta C_{j,t-1}^1 + \epsilon_{i,j,t-1}$$

in the random condition and

$$\text{IV-E1: } 1(C_{i,t-1}^1, C_{i,t-1}^2 = 1) = \alpha_i + \gamma_{t-1} + \beta * 1(\lambda_{t-1} = 0.6) + \delta C_{j,t-1}^1 + \epsilon_{i,j,t-1}$$

in the endogenous condition.  $C_{i,t}^1$  and  $C_{i,t}^2$  are indicators for whether individual  $i$  cooperated in the first and second stage respectively in round  $t$ . The  $C^1$  indicator with a  $j$  subscript represents the first-stage cooperation decision of the partner as a control variable. Individual and round fixed effects are included in both stages. The estimation sample for IV-E1 is limited to those whose partners selected  $\lambda_{t-1}$ . Partner second-stage cooperation in round  $t - 1$  enters as a control in the second stage but not the first because of the timing of the decisions.

Using  $\hat{\epsilon}_{i,j,t-1}$ , the predicted residual from the first stage, the second stage specification in the random condition is

$$\begin{aligned} \text{IV-R2: } 1(C_{i,t}^1, C_{i,t}^2 = 1) = & \alpha_i + \gamma_t + \zeta * 1(C_{i,t-1}^1, C_{i,t-1}^2 = 1) + \eta \hat{\epsilon}_{i,j,t-1} + \\ & \beta_1 \lambda_t + \beta_2 \lambda_t^2 + \beta_3 \lambda_t^3 + \delta_1 C_{k,t}^1 + \delta_2 C_{j,t-1}^2 + \delta_3 C_{j,t-1}^1 + \epsilon_{i,j,k,t} \end{aligned}$$

In the endogenous condition, we introduce an additional control for whether individual  $i$ 's choice of  $\lambda$  is implemented in round  $t$ ,

$L_{i,t}$ . The sample is again restricted to those who did not choose  $\lambda$  in the previous game.

$$\text{IV-E2: } 1(C_{i,t}^1, C_{i,t}^2 = 1) = \alpha_i + \gamma_t + \zeta * 1(C_{i,t-1}^1, C_{i,t-1}^2 = 1) + \eta \hat{\epsilon}_{i,j,t-1} + \\ \beta * 1(\lambda_t = 0.6) + \delta_1 C_{k,t}^1 + \delta_2 C_{j,t-1}^2 + \delta_3 C_{j,t-1}^1 + \theta L_{i,t} + \epsilon_{i,j,k,t} .$$

The  $j$  subscript continues to represent individual  $i$ 's partner in round  $t - 1$  and  $k$  is introduced to represent individual  $i$ 's partner in round  $t$ . This control function approach is implemented in two-stages with standard errors clustered at the individual level in the second stage. Because the standard errors from the manual two-stage procedure fail to account for the estimated nature of the instrument, we also present results using an automated procedure that adjusts the standard errors but does not respect the timing of the control variables between the first and second stages.<sup>18</sup>

The instrument sets are both relevant. The first stage for the random condition yields an  $F(3, 109)$ -statistic of 4.45,  $p < 0.01$  on the joint test of the first, second and third-order  $\lambda$  coefficients being equal to zero. In the endogenous condition, the indicator for  $\lambda = 0.6$  has a positive and significant effect on the likelihood of cooperation in both stages of the game, with  $p < 0.01$ .

Our estimates of reinforcement learning are found in Table 7. OLS estimates of the relationship between lagged cooperation and present cooperation yield similar results in both conditions, and the effect is positive and significant.<sup>19</sup> Instrumenting for lagged cooperation generates a much larger coefficient in the endogenous condition only. This is surprising: the OLS estimates would be biased upwards if time-varying personal factors that led to cooperation in the previous round also led to cooperation in the current round. An advantage of the two-stage approach is that the coefficients on the lagged cooperation indicator are simple to test across models despite the regression specifications being different because the second stage is implemented using OLS. We find that the large difference between the two conditions identified in the two-stage IV model is significant at the 10% confidence level ( $p = 0.09$ ). Given the binary dependent variable, the magnitude of the IV coefficient on lagged cooperation in the endogenous condition needs to be taken in context with the large negative influence of the first stage residual. This indicates that the causal effect of lagged cooperation on present cooperation is partly masked by the endogeneity of past behavior, although this endogeneity is not in the intuitive direction. While the estimates are noisy, the meaningful magnitudes of the coefficients indicate that cooperation is more strongly learned when it arises from an endogenously designed interaction.

## 5 Discussion and Conclusion

People frequently enter into short term relationships where there is much unknown about their partners. An important aspect often under the control of people is how they sequence the values at stake in each interaction. Our intuition suggests, and prior research has shown, that in such cases it would be best to start small—if an interaction goes well, players can feel more comfortable increasing the stakes. In this paper, rather than controlling the relative size of the stakes, we let subjects choose the stake allocation of payoffs across a twice played prisoners' dilemma game. We ask, will they choose to start small, and will they choose the relative stakes that are known to maximize payoffs when assigned exogenously?

Our study builds directly on the work of Andreoni and Samuelson (2006), in which individuals in a laboratory experiment play a series of twice-played prisoners' dilemma games. Their innovation was to experimentally vary the allocation of potential surplus

<sup>18</sup>In other words, variables that should be excluded from the first stage cannot be, using the packaged statistical approach that allows for the proper imputation of standard errors. Manual adjustments of the standard errors in the context of the two-stage models are difficult given that the random and endogenous specifications are estimated simultaneously to allow for hypothesis testing.

<sup>19</sup>Dynamic panel fixed effects models are known to be inconsistent and biased towards zero (Nickell, 1981). However, our goal in this exercise is to compare across conditions rather than interpret the estimate magnitudes themselves.

**Table 7: IV Estimates of Reinforcement Learning in Cooperation**

	OLS Estimates		IV Estimates			
	Random	Endogenous	2 Stage		1 Stage	
			Random	Endogenous	Random	Endogenous
Cooperated Last Round?	0.233*** (0.049)	0.299** (0.114)	0.210 (0.239)	0.906*** (0.332)	0.318 (0.244)	1.084* (0.605)
Last Round Residual			0.023 (0.242)	-0.619* (0.337)		
<i>N</i>	2086	326	2086	326	2086	326

Notes: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . Standard errors, clustered by individual are in parentheses under the estimates. Two stage least squares is implemented using the control function approach and the estimating equations described in the text. The one-stage procedure ignores the timing of variable determination, but has the advantage of factoring first stage noise into the standard errors in the second stage.

across the two stages of each prisoners' dilemma, allowing them to estimate the distribution of stakes that maximized total surplus across both stages. Starting small, with around two-thirds of the potential reserved for the second stage, maximized total social surplus in the game. Here we ask the next, and perhaps more important question: Will subjects gravitate toward starting small? Will choices be concentrated on what Andreoni and Samuelson found to be the surplus maximizing allocation of surplus? Will earnings at this optimum allocation be the same as when the stakes were experimentally controlled?

We find that starting small remains optimal for inducing cooperation and maximizing social surplus when the stakes are chosen by the participants themselves. The payoff maximizing allocation of stakes is nearly identical in the experimental and the endogenous choice of stakes. Additionally, we find evidence of learning to start small over the course of the study. Subjects are significantly more likely to start small, and choose stakes significantly closer to the payoff-maximizing allocation, as the study progresses. Moreover, this result is robust to alternative measures.

We also found an effect that is not directly predicted by Andreoni and Samuelson (2006). When we compare the distributions of earnings conditional on the allocation of stakes, we find that those stakes that are nearer to the payoff maximizing stakes are *more* profitable in the endogenous condition than in the control condition, while those allocations of stakes far away from the optimum earn *less* in the endogenous condition. Stated differently, there is a level effect of the endogenous condition on the expected payoff around its maximum. These excess earnings come from a cooperation-inducing effect of the condition, conditional on stakes, that is largest at the most commonly chosen values. Additionally, we present evidence of what could be interpreted as stronger reinforcement learning of cooperation in the endogenous condition. We speculate that individuals believe they are gaining information about their partner's type through their choice of stakes and that they are more inclined to cooperation when this information is reassuring. The theory posed by Andreoni and Samuelson was silent on the effect of allowing the allocation of stakes to be endogenous. This evidence suggests there may be further gains from allowing people to choose the allocation of stakes themselves, perhaps because it allows more information to be gleaned about one's own character.

This result also speaks more generally to the ingenuity of individuals in structuring their interactions so as to achieve desirable outcomes. This is a topic that economics often leaves unaddressed, but which (as we believe we have shown here) presents much promise. As numerous laboratory and field experiments have shown, many individuals behave pro-socially in social dilemmas, largely based on moral principle or altruistic intention. It is intuitive that people would structure interactions to make the greatest

advantage of such “principled agents,” and doing so is reinforced by the improved payoffs. Looking for natural structures like starting small as evidence of successfully building on the tastes of these principled-agents is one way to pursue this observation. Another is for the economic theorist to use what has been learned about the heterogeneity of tastes in various social dilemmas to identify simple, self-directed measures that society can take to improve the efficiency of such interactions.

## References

- [1] T.K. Ahn, R. Mark Isaac, and Timothy C. Salmon. Endogenous group formation. *Journal of Public Economic Theory*, 10(2):171–194, 2008.
- [2] T.K. Ahn, R. Mark Isaac, and Timothy C. Salmon. Coming and going: Experiments on endogenous group sizes for excludable public goods. *Journal of Public Economics*, 93(1-2):336–351, 2009.
- [3] S. Nageeb Ali and David A. Miller. Enforcing cooperation in networked societies. Working Paper, 2013.
- [4] James Andreoni. Trust, reciprocity and contract enforcement: Experiments on satisfaction guaranteed. Working Paper, 2014.
- [5] James Andreoni and Laura Gee. Gun for hire: Delegated enforcement and peer punishment in public goods provision. *Journal of Public Economics*, 96(11-12):1036–1046, 2012.
- [6] James Andreoni and John Miller. Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70:737–753, 2002.
- [7] James Andreoni and Larry Samuelson. Building rational cooperation. *Journal of Economic Theory*, 127(1):117–154, 2006.
- [8] James Andreoni and Hal Varian. Preplay contracting in the prisoners dilemma. *Proceedings of the National Academy of Sciences*, 96(19):10933–10938, 1999.
- [9] Ken Binmore, Chris Proulx, Larry Samuelson, and Joe Swierzbkiski. Hard bargains and lost opportunities. *The Economic Journal*, 108:1279–1298, 1998.
- [10] Matthias Blonski and Daniel Probst. The emergence of trust. Working Paper, 2004.
- [11] Colin Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Russell Sage Foundation and Princeton University Press, Princeton, 2003.
- [12] Gary Charness and Chun-Lei Yang. Starting small toward voluntary formation of efficient large groups in public goods provision. *Journal of Economic Behavior and Organization*, 102:119–132, 2014.
- [13] Matthias Cinyabuguma, Talbot Page, and Louis Putterman. Cooperation under the threat of expulsion in a public goods experiment. *Journal of Public Economics*, 89(8):1421–1435, 2005.
- [14] Saikat Datta. Building trust. *STICERD Discussion Paper*, 305, 1996.
- [15] Douglas W. Diamond. Reputation acquisition in debt markets. *Journal of Political Economy*, 97(4):828–862, 1989.
- [16] John Duffy, Jack Ochs, and Lise Vesterlund. Giving little by little: Dynamic voluntary contribution games. *Journal of Public Economics*, 91:1708–1730, 2007.
- [17] Karl-Martin Ehrhart and Claudia Keser. Mobility and cooperation: On the run. Working Paper, 1999.
- [18] Ernst Fehr and Simon Gächter. Cooperation and punishment in public goods experiments. *The American Economic Review*, 90(4):980–994, 2000.
- [19] Thomas Fujiwara, Kyle C. Meng, and Tom Vogl. Estimating habit formation in voting. NBER Working Paper No. 19721, 2013.
- [20] Avner Greif and David D. Laitin. A theory of endogenous institutional change. *American Political Science Review*, 98:633–652, 2004.

- [21] Martin Kocher, Peter Martinsson, and Martine Visser. Social background, cooperative behavior and norm enforcement. *Journal of Economic Behavior and Organization*, 81(2):341–354, 2012.
- [22] David M. Kreps, Paul Milgrom, John Roberts, and Robert Wilson. Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, 27(2):245–252, 1982.
- [23] Leslie M. Marx and Steven A. Matthews. Dynamic voluntary contribution to a public project. *Review of Economic Studies*, 67(2):327–358, 2000.
- [24] Nikos Nikiforakis. Punishment and counter-punishment in public good games: Can we really govern ourselves? *Journal of Public Economics*, 92(1-2):91–112, 2008.
- [25] U.S. Department of Justice and the Federal Trade Commission. Horizontal merger guidelines. August 19, 2010.
- [26] Talbot Page, Louis Putterman, and Bulent Unel. Voluntary association in public goods experiments; reciprocity, mimicry and efficiency. *Economic Journal*, 115(506):1032–1053, 2005.
- [27] Michael Peters. Decentralized markets and endogenous institutions. *The Canadian Journal of Economics*, 28(2):227–260, 1995.
- [28] David G. Rand, Anna Dreber, Tore Ellingsen, Drew Fudenberg, and Martin A. Nowak. Positive interactions promote public cooperation. *Science*, 325(5945):1272–1275, 2009.
- [29] James E. Rauch and Joel Watson. Starting small in an unfamiliar environment. *International Journal of Industrial Organization*, 21(7):1021–1042, 2003.
- [30] Thomas C. Schelling. *Strategy of Conflict*. Harvard University Press, 1960.
- [31] Joel Watson. Starting small and renegotiation. *Journal of Economic Theory*, 85(1):52–90, 1999.
- [32] Joel Watson. Starting small and commitment. *Games and Economic Behavior*, 38(1):176–199, 2002.
- [33] Roberto A. Weber. Managing growth to achieve efficient coordination in large groups. *The American Economic Review*, 96(1):114–126, 2006.