Building Rational Cooperation On Their Own: Learning to Start Small

James Andreoni[*]

University of California, San Diego
& NBER

Michael A. Kuhn[†]

University of Oregon

Larry Samuelson[‡]

Yale University

March 22, 2018

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 willing to 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 gravitate toward the payoff maximizing strategy of starting small in a twice-played prisoners' dilemma. Intriguingly, the salutary payoff effects of doing so are larger than those that arise when the same allocation is exogenously chosen.

JEL Classification: C92, D64, Z13
Keywords: Cooperation, Starting Small, Learning, Prisoners' Dilemma

1 Introduction

Can groups of individuals informally develop institutions and structure interactions that take advantage of the naturally occurring dispositions of many people to prefer cooperation? And can this happen even in interactions that are too short to build reputations?

We address these questions with a laboratory experiment that builds on the theoretical framework and experimental work of Andreoni and Samuelson (2006). These authors examined a class of twice-played "prisoners' dilemmas," with two distinguishing characteristics. First, while the total monetary amount at stake over the two periods is fixed, different versions of the game distribute these stakes across the two periods differently. It may be that the two iterations of the prisoners' dilemma are played for the same stakes, as is customary, but may also be that the stakes are larger in the first period than the second or vice versa.

Second, reflecting the quotation marks in the previous paragraph, we assume the players have preferences that cause the utilities from cooperating and defecting to increase in the probability that their opponent cooperates (as usual), but with the utility from cooperating increasing faster, to the extent that cooperation may yield a higher utility than defection, *if* the opponent is sufficiently likely to cooperate. Players are heterogeneous, differing in the likelihood of opponent cooperation required to ensure that the utility of cooperation exceeds that of defection.

Andreoni and Samuelson (2006) identify conditions under which the equilibrium joint payoffs in such a setting are maximized if the game "starts small," so that the second-period stage game is played for higher stakes than those of the first period. Their experimental results confirm this intuitive hypothesis. Joint payoffs are maximized by playing for approximately one third of the total stakes in the first stage, reserving two-thirds for the second stage.

The arrangement of stakes across the periods in Andreoni and Samuelson (2006) is fixed exogenously. In this paper we reproduce the Andreoni-Samuelson experimental game, but this time we allow the subjects themselves to determine the relative stakes. We find that joint payoffs are maximized by an arrangement of stakes nearly identical to that found by Andreoni and Samuelson. Moreover, the subjects indeed gravitate toward this same allocation of stakes.

We are especially intrigued by the finding that the gains from arriving at this allocation are significantly higher when the relative stakes are chosen by the players rather than controlled experimentally. That is, when the subjects choose to start small on their own, it generates more cooperation than when those same stakes are set by the experimenter. This result opens up new questions for theorists, experimenters, and policy makers. In particular, how well and how often can decentralized groups of people endogenously learn and develop structured ways of interacting that help them collectively achieve more efficient outcomes? And does mutual recognition of the strategic sophistication of partners aid in this development?

The next section provides some background on starting small and the endogenous determination of relationship stakes. We describe our experimental procedures in section 3, present the results in section 4, and conclude in section 5.

2   Background on Starting Small

In this section we briefly discuss the literature on starting small and rational cooperation. We then provide an intuitive description of the theoretical model presented in Andreoni and Samuelson (2006), and finally discuss how this theory could generalize to a game where the size of the stakes is chosen by the players.

2.1   The Literature

We build on four strands of literature. First, the underlying theoretical model presented in Andreoni and Samuelson (2006) is a finitely repeated game of incomplete information. Kreps, Milgrom, Roberts and Wilson (1982) highlighted the role of incomplete information in the finitely-repeated prisoners' dilemma, giving rise to a flourishing literature summarized in Mailath and Samuelson (2006, chapter 17). We differ from much of this literature in focusing on short (two-period) games.

Second, the finitely-repeated-games literature emphasizes that even minuscule amounts of heterogeneity in agents' preferences can have significant effects on equilibrium play (if the game is sufficiently long). We join the extensive literature on social preferences in thinking that people whose preferences are based on more than simply monetary payoffs 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.[1]

Third, our work most directly fits into a small but growing literature examining the virtues of starting small. Schelling (1960) suggests an incremental approach to funding public goods, an idea formalized by Marx and Matthews (2000) and examined experimentally by Duffy, Ochs, and Vesterlund (2007). Sobel (1985), Ghosh and Ray (1996) and Kranton (1996) examine various settings in which equilibrium strategies embed a notion of starting small: a credibility-building game between a lender and borrower, communities seeking to achieve cooperation with limited information about past behavior, and partnerships formed and maintained in the constant presence of the outside option to start over with someone else, respectively. Closely related in context to our study, 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

---

[1]See Andreoni and Miller (2002) and Camerer (2003, chapter 2).

small plays a role in developing commercial relationships in developing countries. These papers often use starting small as a means to increase the effective discount rate, in contrast to our focus on short relationships in which discounting plays no meaningful role.[2] Kamijo, Ozono and Shimizu (2016) and Ozono, Kamijo and Shimizu (2016) report experiments in which increasing stakes can facilitate coordination in coordination games and public goods games.

Fourth, the literature includes some similarly motivated studies in which the players choose some aspect of the game they are to play. The literature on punishment in public goods games indicates that exogenously engineered opportunities to punish can be destructive, while the endogenous adoption of delegated enforcement can be more effective.[3] Related work by Andreoni (2017) 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. Peters (1995) 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. 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, values that are fixed in the short run, but variable in the long run. Other work has examined the endogenous determination of the players in the game (rather than the specification of the game). Charness and Yang (2014) use the laboratory to investigate how behavior, earnings, and efficiency can be enhanced by a voting procedure that allows groups participating in a public-goods game to determine their own members.[4] Ali and Miller (2013) study a theoretical model of a networked society in which the formation of each link is endogenously determined by individuals. Altogether, one could view this recent work (as well as the present paper) as creating a framework for studying the development of relationships, communities and enforcement mechanisms in an environment otherwise devoid of institutions.

2.2 Theoretical Intuition from Andreoni and Samuelson (2006)

This section provides an informal discussion of the model and results of Andreoni and Samuelson (2006), counting on readers to refer to the original for details. Two players play a "prisoners'

---

[2]Others that develop theoretical models in which starting small optimally builds relationships include 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.

[3]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).

[4]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).

Figure 1: Two-Stage Prisoners' Dilemma, $0 \leq x_1 \leq 10$ and $x_1 + x_2 = 10$.

dilemma," observe the outcome, and then (without discounting) play the stage game again. Figure 1 presents the parameters of the games used here and by Andreoni and Samuelson (2006).

The variables x_1 and x_2 determine the stakes for which the game is played in each stage, with $0 \leq x_1 \leq 10$ 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 λ is the fraction of total payoffs reserved for stage 2. Starting small means $\lambda > 1/2$.

The players in the model are heterogeneous. We suppose that each player's preferences can be characterized by a number α, where an individual playing a single prisoners' dilemma will prefer to cooperate *if* they believe their opponent will cooperate with a probability at least α. We say those with lower values of α are "more altruistic." The values of α range from below 0 (in which case the player always cooperates) to above 1 (always defect). In a single prisoners' dilemma, there would be at least one fixed point where exactly α^* fraction of the population have preference parameters less than or equal to α^*, and there would exist a corresponding equilibrium in which proportion α^* of the players cooperate.

To build intuition for the twice-played prisoners' dilemma, 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 α'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 α. This gives rise to a critical value α_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. On one hand, it increases the desire to pool with lower-type α'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

4

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 α types, tending to decrease cooperation in the second stage. If the distribution of α is smooth, then when we make a small movement away from equal stakes toward larger stakes in stage 2, we can reasonably expect the first effect to 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.

2.3 Generalizing to Endogenous Stakes

The setting examined in this paper differs from the Andreoni-Samuelson 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 every player, regardless of their cooperative intent or type, selects the expected payoff maximizing allocation of stakes and duplicates the play found in Andreoni and Samuelson (2006), conditional on having such stakes exogenously set.

In light of this, we investigate three questions. First, are joint payoffs maximized at the same value λ^* as in Andreoni and Samuelson (2006)? Second, there is no reason to believe that all subjects have equal or accurate priors on the distributions of preferences in the sample, nor do experimental subjects typically immediately hit on equilibrium play. We accordingly ask, do the stakes chosen by the subjects gravitate toward λ^* over the course of the experiment? Third, does the subjects' behavior and the corresponding payoffs, for endogenously chosen stakes near λ^*, duplicate those found in Andreoni and Samuelson (2006)?

3 Experimental Procedures

We examine data from a total of eight experimental sessions, including five from the original Andreoni-Samuelson paper, where λ is chosen by the experimenter, and three new sessions where

λ is chosen by the players themselves.[5] 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 λ. We will call this the *short sample*. Given our interest in subjects learning over time, 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. 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 λ is endogenously determined by subjects, and a random condition, referring to the original data in which the computer randomly drew λ from a discrete distribution ranging from zero to one with equal weight on each 0.1 increment, including both ends. The original data come from 5 sessions each involving 22 subjects playing 20 rounds of the twice repeated prisoners' dilemma, implying 1100 interactions. For all side-by-side comparisons of the original and new data that follow, we exclude the new 40 round session.[6]

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 λ 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 λ chosen, and not which player selected the value.[7] Subjects were paid for their performance in all games in cash following the experiment.

---

[5]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.

[6]Results from this long sample are very similar to those from the short samples, so we only present results from the long sample when we wish to focus on issues specific to the experiment length.

[7]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 λ is used will know the value of λ chosen by the other player. This fact is a constant across all games.

4 Results

We present our results in three parts. In subsection 4.1 we first ask whether the λ 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 they are nearly identical. Next, in subsection 4.2, we present evidence that subjects are indeed migrating towards the joint-payoff maximizing value of λ. Third, we show in subsection 4.3 that payoffs, conditional on λ, differ under the random and endogenous conditions, with higher payoffs appearing in the endogenous conditions, especially at the values of λ close to the optimum.

4.1 What Value of λ Maximizes Joint Payoffs?

Figure 2 presents the mean joint payoffs from a single (two-stage) interaction separately by λ and by condition. As can easily be seen, $\lambda = 0.6$ provides the maximum payoff for *both* the random and endogenous conditions, indicating that selecting λ endogenously did not overturn the result based on an exogenously chosen λ.
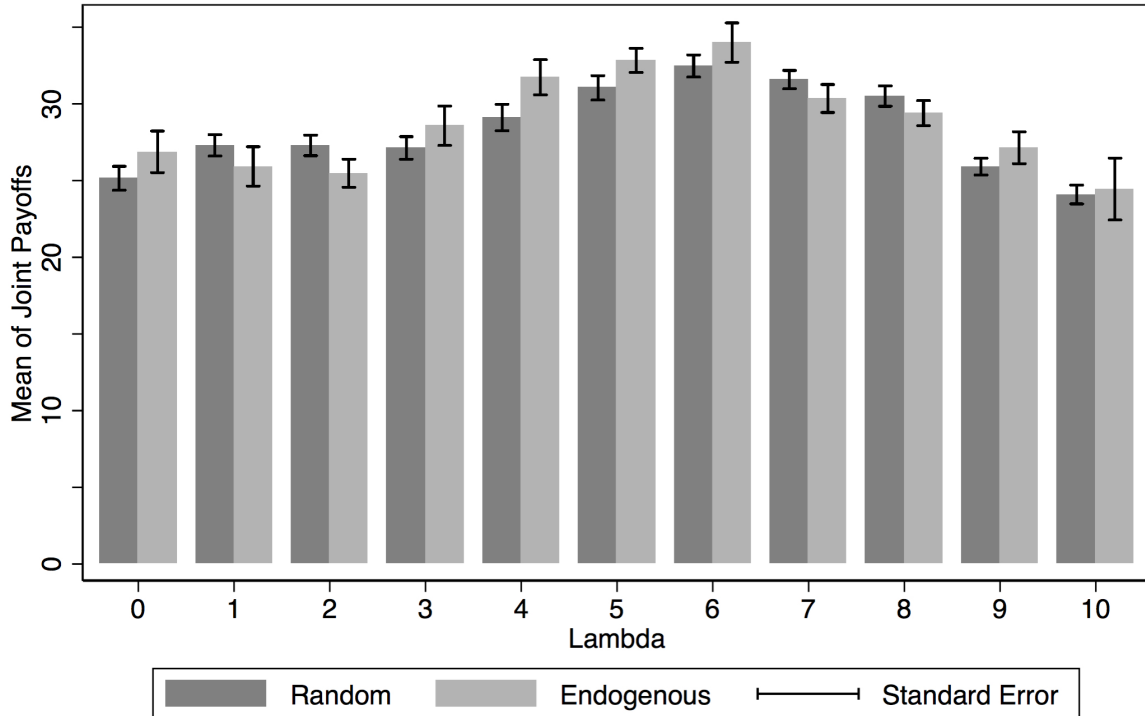

Figure 2: Payoffs by Value of λ and Condition

To establish the statistical significance of this observation, we follow Andreoni and Samuelson (2006) and estimate joint payoffs, π, as a cubic polynomial of λ, conditional on a round fixed

7

effect, γ_t. We then find the value of λ that maximizes this polynomial.[8] Of course, individual characteristics may play a role in determining the chosen λ in the endogenous condition, and these same personal characteristics are likely to influence how people play the game once λ is determined, and thus how much they earn from playing. To account for this, we augment the specification with individual fixed effects:

$$\pi_{i,t} + \pi_{j,t} = \theta_i + \theta_j + \gamma_t + \beta_1 \lambda_{k(i,j),t} + \beta_2 \lambda_{k(i,j),t}^2 + \beta_3 \lambda_{k(i,j),t}^3 + \epsilon_{i,j,t} \, , \tag{1}$$

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 λ is chosen. θ_i and θ_j are individual-specific constants for both players in a pairing. An important consideration for standard errors is that the unit of analysis is the game that features a unique pairing of subjects. We apply two-way clustering using each individual within a pair.[9]

   We estimate equation (1) allowing β coefficients to differ across conditions. Results are reported in Table 1 for the full sample, and samples limited to the first 10 rounds and last 10 rounds. We never reject the null hypotheses that set of β coefficients or the payoff-maximizing λ are the same across conditions. Restricting attention to the last 10 rounds, we see a slight drop in the payoff-maximizing value λ^* from 0.673 in the random condition to 0.623 in the endogenous condition, but the two estimates of λ^* are not significantly different at conventional levels.[10]

4.2   What Value of λ do Players Choose?

Here we first ask whether subjects see and learn the strategy of starting small. We then look more specifically at the λ^* 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. The evolution in choices is also visible in the changing empirical distribution of λ over time. We present histograms of λ choices for the early and late rounds in Appendix Section B, along with statistical tests of the difference in distributions.

---

[8]Estimating a cubic allows the data more flexibility to contradict the predictions of this relationship from Andreoni and Samuelson (2006). In practice, we do observe a hump-shaped relationship as hypothesized, and the cubic gives us the flexibility to estimate diminishing sensitivity of payoffs to λ in the start-large domain.

[9]Subjects in our study are randomly and anonymously rematched each round. Thus, the potential exists for the error term across pairings within a round to be correlated because of shared past experiences. There is no perfect way to implement clustering within a session. We can be more conservative, but at the cost of being ad-hoc—for example, manually re-scaling the standard errors to reflect a sample size equal to the number of individuals in the sample—and obtain similar results.

[10]We also estimate equation (1) without individual specific fixed effects, comparing the difference for the endogenous condition only. Relative to a model with fixed effects we find significantly different coefficients on the cubic polynomial estimates, yet nearly identical values of λ^* with and without the individual fixed effects, and those values are also nearly identical to those reported in Table 1. All measures of λ^* are not significantly different. This can be seen in appendix, Section A, Table A1.

Table 1: Relationship between λ and Payoffs across Conditions

	Sample Restriction		
	All Rounds	Rounds 1-10	Rounds 11-20
Random Condition Terms:			
λ	-1.408	-12.623	-3.205
	(6.693)	(11.466)	(6.949)
λ^2	43.823***	66.310**	49.563***
	(15.883)	(26.457)	(15.088)
λ^3	-44.768***	-59.212***	-46.765***
	(10.687)	(17.495)	(9.463)
Endogenous Interactions:			
$\lambda * 1(endogenous = 1)$	4.186	19.080	5.619
	(12.847)	(24.381)	(16.906)
$\lambda^2 * 1(endogenous = 1)$	-6.602	-57.092	2.515
	(31.222)	(61.421)	(38.373)
$\lambda^3 * 1(endogenous = 1)$	4.819	46.462	-10.990
	(21.943)	(42.312)	(25.209)
H_0: Endog. Interactions = 0	$F(3, 1361) = 0.28$	$F(3, 601) = 1.07$	$F(3.601) = 1.09$
	$(p = 0.84)$	$(p = 0.36)$	$(p = 0.35)$
Random Payoff-max. λ_r^*	0.636	0.635	0.673
	(0.020)	(0.029)	(0.019)
Endogenous Payoff-max. λ_e^*	0.656	0.717	0.623
	(0.030)	(0.178)	(0.031)
$H_0 : \lambda_r^* = \lambda_e^*$	$\chi^2(1) = 0.31$	$\chi^2(1) = 0.21$	$\chi^2(1) = 1.76$
	$(p = 0.58)$	$(p = 0.65)$	$(p = 0.19)$
N	1540	770	770

Notes: $^{**}p < 0.05$, $^{***}p < 0.01$. Standard errors with two-way clustering for both individuals in a pairing are in parentheses under the estimates unless otherwise indicated. Round fixed-effects and individual fixed-effects are included in all specifications. The payoff-maximizing λ is a non-linear combination of the three coefficient estimates obtained using the quadratic formula on the derivative of the implied cubic polynomial. Standard errors for the payoff-maximizing λ are calculated via the delta method.

We partition our sample into three intervals. Call rounds 1 to 6 the beginning, rounds 7-14 the middle, and rounds 15-20 the end. We sort subjects based on their choice of λ. 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 2, 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 small at the beginning do so to even a greater degree.[11]

Table 2: Evolution of Sub-Sample Sizes over Time

Group, Rounds	Fraction of Sample in Each Sub-Sample		
	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			

Would subjects have continued to learn to start small after round 20? We turn to the long sample to answer this. To do the same analysis in the 40-round sample, we use intervals of the same

---

[11]Paired t-tests of the short sample frequencies in Table 2, 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.

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 Panel B of Table 2, 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 magnitude of the shift towards starting small is larger in the long sample, it is worth noting that starting large is less prominent overall in the long sample.[12]

The distributions of λ choices that we present in Appendix Section B tell a similar story. In the short sample, there is a clear shift of mass from the left side of the distribution (starting large) to the right ride of the distribution (starting small) over time. The same is true in the long sample, however, whereas starting even is still relatively common at the end of the short sample, other values in the start small region overtake it in frequency by the end of the long sample.

Table 3: Time Trends in λ Choice

	Dependent Variable				
	(1)	(2)	(3)	(4)	(5)
	$\lambda_{i,t}$	$\lambda_{i,t} - \lambda_{i,t-1}$	$\|\lambda_{i,t} - 0.656\|$	$Pr(\lambda_{i,t} \in [0.6, 0.9])$	H-index
	Panel A: Short Sample				
Round	0.005***	0.002**	-0.003**	0.007	0.002**
	(0.002)	(0.001)	(0.001)	(0.004)	(0.001)
Constant	0.456	-0.018	0.265	0.306	0.107
	(0.029)	(0.010)	(0.018)	(0.053)	(0.007)
N	787	705	787	880	20
	Panel B: Long Sample				
Round	0.003**	0.001***	-0.001*	0.007**	-0.000
	(0.001)	(0.000)	(0.000)	(0.003)	(0.000)
Constant	0.499	-0.016	0.214	0.350	0.149
	(0.027)	(0.007)	(0.010)	(0.061)	(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.

---

[12] Paired t-tests of the long sample frequencies in Table 2, 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 ($p = 0.10$).

In Table 3 we look at the changes in λ throughout the sessions from a number of vantage points. In column (1) we start with λ itself, to ascertain an overall directional trend. In column (2) we then estimate a first-differenced specification. Next, in column (3) we consider the absolute deviation of λ from 0.656, our estimate of the payoff-maximizing λ 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 λ.[13] The last two specifications do not require this exclusion. As an analogue to Table 2, 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 concentrated the market for λ values is. This approach is designed to assess whether learning and convergence happen over the course of a session.[14]

The regression results shown in Table 3 include standard errors clustered by individual in all specifications except for the H-index specification. We find, in both the short and long samples, that the mean choice of λ is increasing slowly and significantly over time. In the short sample, the predicted λ 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, the coefficients on the round variable shown in Table 3 indicates that subjects are growing closer to choosing the payoff maximizing λ as play continues, and these coefficients are precisely measured in 4 of 5 tests for both the short and long samples.

4.3 Do Players Earn More When λ is Endogenous?

Figure 3 presents the difference between the mean joint payoffs in the random condition and the mean joint payoffs in the endogenous condition, for each round. As the sessions proceed, this difference grows. This is expected. In the random condition, the (randomly chosen) value of λ is often quite far from its optimal value of λ^*. In the endogenous condition, the subjects' choices tend toward the optimal value λ^*, allowing them to achieve higher payoffs. Table 4 shows regressions of joint payoffs on a dummy variable for the endogenous condition. Columns (1) and (4) shows the treatment effect of the endogenous condition relative to the random condition. Payoffs are about 10% larger over the last 10 rounds, an effect which is statistically significant at a 5% level.

---

[13]Not only are these extreme values cardinally ambiguous with respect to the other λ values, but with respect to each other as well. $\lambda = 0$ corresponds to a single-shot game for all the payoffs followed by a meaningless game. $\lambda = 1$ corresponds to a cheap talk game followed by a single-shot game for all the payoffs. Moving from $\lambda = 0.9$ to $\lambda = 1$ ($\lambda = 0.1$ to $\lambda = 0$) may therefore be very different than moving from $\lambda = 0.8$ to $\lambda = 0.9$ ($\lambda = 0.2$ to $\lambda = 0.1$).

[14]Given our 10 distinct values of λ (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 λs 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.
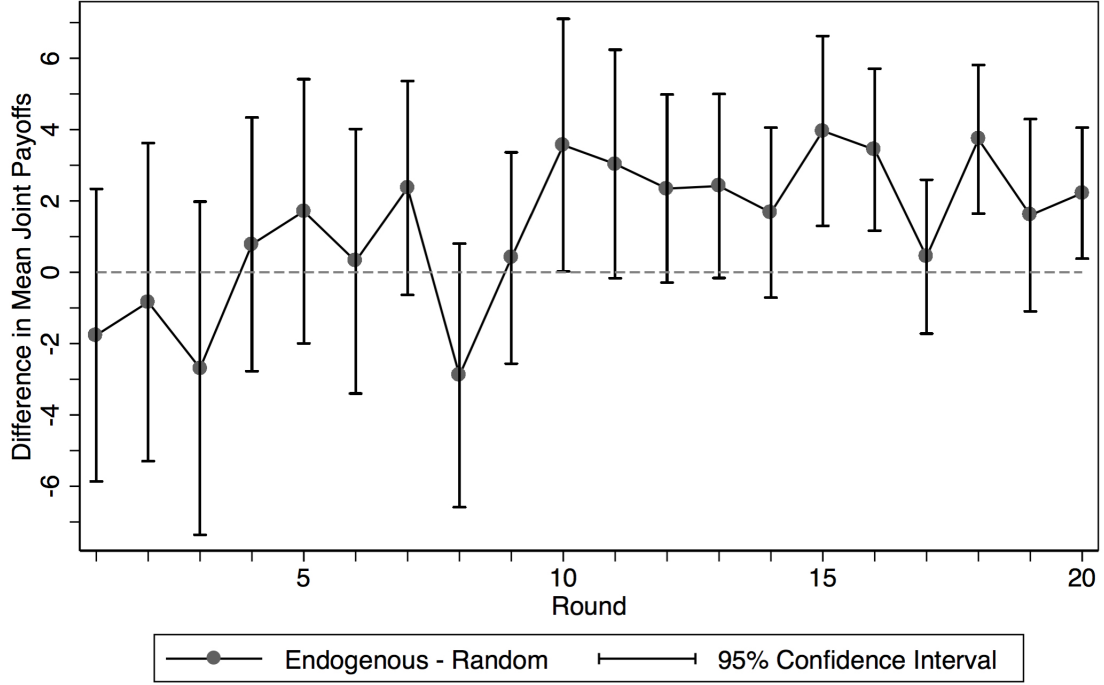
Figure 3: Excess Payoffs in Endogenous Condition over Random Condition by Round

The excess payoffs in the endogenous condition seen in Figures 2 and 3 lead us to examine whether there are payoff differences *conditional* on λ. In other words, once λ is chosen, does it matter if that choice came from a computer or a player? These estimates represent the combination of two effects: 1) the causal effect of a value of λ being endogenously rather than randomly derived, and 2) the effect of selection into a value of λ. Because these effects may differ considerably depending on what λ is, we define two groups of λ choices. We call λ values of 0.4, 0.5, or 0.6 *central* choices. These choices reveal that a player would like to split the payoffs roughly evenly across the two stages. All other values of λ we call *extreme* values. Columns (2) and (5) of Table 4 show regressions of joint payoffs on a dummy variable for the endogenous condition, limited to central λ values only.[15] We find that payoffs at central values of λ are about 19% larger in the second half of the endogenous condition sessions relative to the second half of the random condition sessions. This is evidence that selection into a well-balanced version of the game (and perhaps the signaling of that intent) can stimulate even higher cooperation than random assignment to the same game. Columns (3) and (6) of Table 4 show regressions of joint payoffs on a dummy variable for the endogenous condition, limited to extreme λ values only. We find a negative, but statistically insignificant impact of selection into an imbalanced version of the game.

If there is signaling value to the selected λ, then values of λ far outside the optimal range

---

[15]We also include λ fixed effects to hold fixed the causal impact of *different* λ values on payoffs.

Table 4: Payoff Differences across Conditions

	Sample Restriction					
	All Rounds			Rounds 11-20		
	(1)	(2)	(3)	(4)	(5)	(6)
$1(endogenous = 1)$	1.290	1.952	-0.233	2.485**	4.396**	-0.400
	(1.190)	(1.679)	(0.947)	(1.142)	(1.990)	(0.790)
Constant	28.374	29.335	25.620	24.442	23.663	21.073
	(0.553)	(1.130)	(1.171)	(0.528)	(1.120)	(0.715)
λ Fixed Effects?	No	Yes	Yes	No	Yes	Yes
λ Range	All	Central	Extreme	All	Central	Extreme
N	1540	446	1074	770	227	543

Notes: $^{**}p < 0.05$. Standard errors with two-way clustering for both individuals in a pairing are in parentheses under the estimates. In columns (2) and (5), $\lambda = 0.04$ is the excluded group for the fixed effects. In columns (3) and (6), $\lambda = 0$ is the excluded group for the fixed effects.

could be warning signs, and we might expect to see a negative effect of the endogenous condition, while, by contrast, we might expect that endogenous values near the optimum could induce more cooperation. The most direct test of whether the choice of λ is interpreted as a signal is to leverage the fact that when a subject's choice of λ is implemented, this news is relatively uninformative. For example, if a subject chooses $\lambda = 0.6$ and then observes $\lambda = 0.6$ implemented, the interpretation should be that there is a greater likelihood that the subject's own choice was implemented than the choice of their partner. However, when a subject chooses $\lambda = 0.6$ and then observes $\lambda = 0.5$ implemented, the subject knows both players' choices of λ and, moreover, can infer that both players have similar strategic intentions. The same is true if two extreme choices of λ are revealed.

Our signaling hypothesis is that when both subjects in a pair choose central values of λ, a subject should be more likely to cooperate when their choice of λ is *not* selected because the revelation of λ is *more* informative. Conversely when both subjects in a pair choose extreme values of λ (by which we mean all non-central values), a subject should be *less* likely to cooperate when their choice is *not* selected for the similar reasons.

To test this hypothesis we regress an indicator for whether an individual cooperates in the first stage of the game on four mutually exclusive indicators: 1) *Central/Used* $= 1$ if an individual's central choice of λ was implemented. 2) *Central/Unused* $= 1$ if an individual's central choice of λ was not implemented and a different central λ was implemented. 3) *Extreme/Used* $= 1$ when an individual's extreme choice of λ was implemented. 4) *Extreme/Unused* $= 1$ when an individual's extreme choice of λ was not implemented and a different extreme λ was implemented. 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 λ choice being misaligned with theirs), we expect that the sign of the *Central/Unused* coefficient to be positive and that on *Extreme/Unused* to be negative. We use a fixed effects for λ, individual fixed effects, and dummy variables for game as controls.

Table 5: Effect of λ on Cooperation in the Endogenous Condition

	First Stage Cooperation	
	All Rounds	Rounds 11-20
Central/Used	0.024	0.112
	(0.079)	(0.102)
Central/Unused	0.021	0.333**
	(0.090)	(0.124)
Difference: *Used−Unused*	0.002	-0.221*
	(0.068)	(0.125)
Extreme/Used	-0.012	0.012
	(0.043)	(0.070)
Extreme/Unused	-0.122**	-0.103
	(0.050)	(0.076)
Difference: *Used−Unused*	0.110***	0.115*
	(0.037)	(0.063)
N	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 λ fixed effects and individual fixed effects.

Results are presented in Table 5 for both the full sample and the final 10 rounds. The predictions are first that coefficients on the dummy variables *Central/Unused* should be positive, while those for *Extreme/Unused* should be negative. This prediction is met in all four cases, with one coefficient reaching statistical significance for each time period. The second prediction is that the difference between subjects with used and unused λs will be negative for central values and positive for extreme values. As the table shows, three of the four estimated differences have the correct sign and reach statistical significance at conventional levels, while one difference (central and all rounds) is a precisely estimated zero. Overall, this analysis is supportive of the hypothesis that people recognize and respond to a signal in the level of λ in their partners.[16]

---

[16]This finding is strengthened by more complete analysis presented in Appendix Section C. Here we show evidence that cooperation has a stronger reinforcement effect in the endogenous condition.

5 Discussion and Conclusion

People frequently enter into short term relationships where much is unknown about their partners. An important aspect often under the control of people is how they sequence the values at stake in each interaction. Common intuition suggests it would be best to start small—if an interaction goes well, players can feel more comfortable increasing the stakes. Andreoni and Samuelson (2006) examined this intuition in a theoretical model of a twice-played Prisoner's Dilemma, and validated this prediction with an experiment. Both the theory and experiment, however, were predicated on the choice of stakes *not* being selected by the players themselves. It would appear both more interesting and more valuable to see that the same conclusions—or stronger—hold when starting small is determined endogenously. This paper squares this circle.

Andreoni and Samuelson's innovation was to experimentally vary the allocation of stakes across the two stages of Prisoners' Dilemma. This allowed them to estimate the distribution of stakes that maximized total surplus. 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 natural and more important question: When subject choose the stakes themselves, will they gravitate toward starting small? If so, will the surplus maximizing allocation of stakes be the same? If they are, then will earnings at this optimum allocation be the same as when the stakes were experimentally controlled?

We find that starting small remains optimal; the payoff maximizing allocation of stakes in our experiment are virtually the same as when stakes are experimentally selected. Additionally, we find evidence of learning to start small over the course of the study. Subjects are significantly more likely to start small, and to robustly choose stakes significantly closer to the payoff-maximizing allocation as the study progresses.

We also found an unpredicted but very interesting effect. When the stakes are nearer to the payoff maximizing stakes they are *more* profitable when selected by subjects than when selected experimentally. Stated differently, there appears to be a signaling value to the level of stakes chosen that heightens the returns to starting small and decreases the returns to starting (very) big. Our speculation, which could be of great interest for further development, is that individuals are gaining information about their partner's character through their choice of stakes, despite the existence of an equilibrium in which players of all types select the same stakes.

This result also speaks more generally to the ingenuity of individuals in structuring their interactions. Rational cooperation is possible in a twice played Prisoners' Dilemma game, but is only possible if there truly are those who are willing to cooperate with sufficient assurance of cooperation from their partners. As numerous laboratory and field experiments have shown, many individuals behave pro-socially in social dilemmas, largely based on moral principles or altruistic

16

intentions. It is intuitive that individuals or groups within society would structure interactions to take the greatest advantage of such "principled agents," especially when doing so is reinforced by the improved payoffs. This suggests a potentially valuable area for research. Can we find natural, organic structures like starting small as evidence that people, on their own, can successfully innovate institutions and rules of interaction that leverage these moral or altruistic preferences for the greater good? These structures need no central planner, no clever mechanism designer, and no external enforcer. Instead, as in this study, informal arrangements are efficiency enhancing because of the existence of a (perhaps very small) well of benevolent individuals.

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. Satisfaction guaranteed: When moral preferences meet moral hazard. Working Paper, 2017.

[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] Parikshit Ghosh and Debraj Ray. Cooperation in community interaction without information flows. *Review of Economic Studies*, 63(3):491–519, 1996.

[21] Avner Greif and David D. Laitin. A theory of endogenous institutional change. *American Political Science Review*, 98:633–652, 2004.

[22] John C. Ham, John H. Kagel, and Steven F. Lehrer. Randomization, endogeneity and laboratory experiments: the role of cash balances in private value auctions. *Journal of Econometrics*, 125:175–205, 2005.

[23] Yoshio Kamijo, Hiroki Ozono, and Kazumi Shimizu. Overcoming coodination failure using a mechanism based on gradualism and endogeneity. *Experimental Economics*, 19(1):202–217, 2016.

[24] 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.

[25] Rachel E. Kranton. The formation of cooperative relationships. *Journal of Law, Economics and Organization*, 12(1):214–233, 1996.

[26] 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.

[27] George J. Mailath and Larry Samuelson. *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, Oxford, 2006.

[28] Leslie M. Marx and Steven A. Matthews. Dynamic voluntary contribution to a public project. *Review of Economic Studies*, 67(2):327–358, 2000.

[29] 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.

[30] U.S. Department of Justice and the Federal Trade Commission. Horizontal merger guidelines. August 19, 2010.

[31] Hiroki Ozono, Yoshio Kamijo, and Kazumi Shimizu. Institutionalize reciprocity to overcome the public goods provision problem. *PLoS ONE*, 11(6), 2016.

[32] Talbot Page, Louis Putterman, and Bulent Unel. Voluntary association in public goods experiments; reciprocity, mimicry and efficiency. *Economic Journal*, 115(506):1032–1053, 2005.

[33] Michael Peters. Decentralized markets and endogenous institutions. *The Canadian Journal of Economics*, 28(2):227–260, 1995.

[34] David G. Rand, Anna Dreber, Tore Ellingsen, Drew Fudenberg, and Martin A. Nowak. Positive interactions promote public cooperation. *Science*, 325(5945):1272–1275, 2009.

[35] James E. Rauch and Joel Watson. Starting small in an unfamiliar environment. *International Journal of Industrial Organization*, 21(7):1021–1042, 2003.

[36] Thomas C. Schelling. *Strategy of Conflict*. Harvard University Press, 1960.

[37] Joel Sobel. A theory of credibility. *Review of Economic Studies*, 52(4):557–573, 1985.

[38] Joel Watson. Starting small and renegotiation. *Journal of Economic Theory*, 85(1):52–90, 1999.

[39] Joel Watson. Starting small and commitment. *Games and Economic Behavior*, 38(1):176–199, 2002.

[40] Roberto A. Weber. Managing growth to achieve efficient coordination in large groups. *The American Economic Review*, 96(1):114–126, 2006.

Appendix

For online publication only. To accompany
James Andreoni, Michael A. Kuhn and Larry Samuelson,
"Starting Small: Endogenous Stakes and Rational Cooperation."

A   Alternative Estimates of the Cubic Polynomial

Following Andreoni and Samuelson (2006), we estimate joint payoffs, π, as a cubic polynomial of λ, conditional on a round fixed effect, γ_t. We then find the value of λ that maximizes this polynomial. We call this first specification CP, for cubic polynomial:

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

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 λ is chosen.

Individual characteristics may play a role in determining the chosen λ, and these same personal characteristics are likely to influence how people play the game once λ is determined and thus how much they earn from playing. To account for this, we take an individual fixed effect approach in specification FE:

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

where θ_i and θ_j are individual-specific constants for both players in a pairing. Given both the individual and round constants, any remaining confounding endogeneity of λ must be within-individual, time-varying covariance between the choice of λ and the cooperation decision.

Table A1 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 λ coefficients yields very similar estimates λ^* of the payoff-maximizing λ, and we do not reject equality of the payoff maximizing λ across specifications. While we use multi-way clustering for the main analysis of game-level outcomes, we use unclustered robust standard errors here to accommodate the simultaneous estimation of both models.

Table A1: Relationship between λ and Payoffs in the Endogenous Condition

	Model	
	CP	FE
λ	3.802	1.487
	(13.289)	(12.166)
λ^2	51.013	40.451
	(33.097)	(29.908)
λ^3	-57.439**	-42.722**
	(22.906)	(20.550)
H_0: CP Terms Jointly = FE Terms	$\chi^2(3)$=7.38*	
	$(p = 0.06)$	
Payoff-maximizing λ: λ^*	0.627	0.649
	(0.025)	(0.030)
$H_0 : \lambda^*_{CP} = \lambda^*_{FE}$	$\chi^2(1) = 0.97$	
	$(p = 0.33)$	
N	440	440

$^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Robust standard errors are reported in parentheses under the estimates unless otherwise indicated. The payoff-maximizing λ is a non-linear combination of the three coefficient estimates obtained using the quadratic formula on the derivative of the implied cubic polynomial.

B Choices of λ Trend Upward

This section elaborates on the results of Section 4.2 by examining the full distribution of choices over the course of the short and long samples. Figure A1 shows histograms of λ choices for the first and last 5 rounds in the short sample, while Figure A2 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 ride 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 A1 and A2 differ from one another.[17] In the short sample, the distribution of λ 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

---

[17]Aggregated up to 5-round bins, testing the distributions against the hypothesis of uniformity rejects the null in all circumstances.
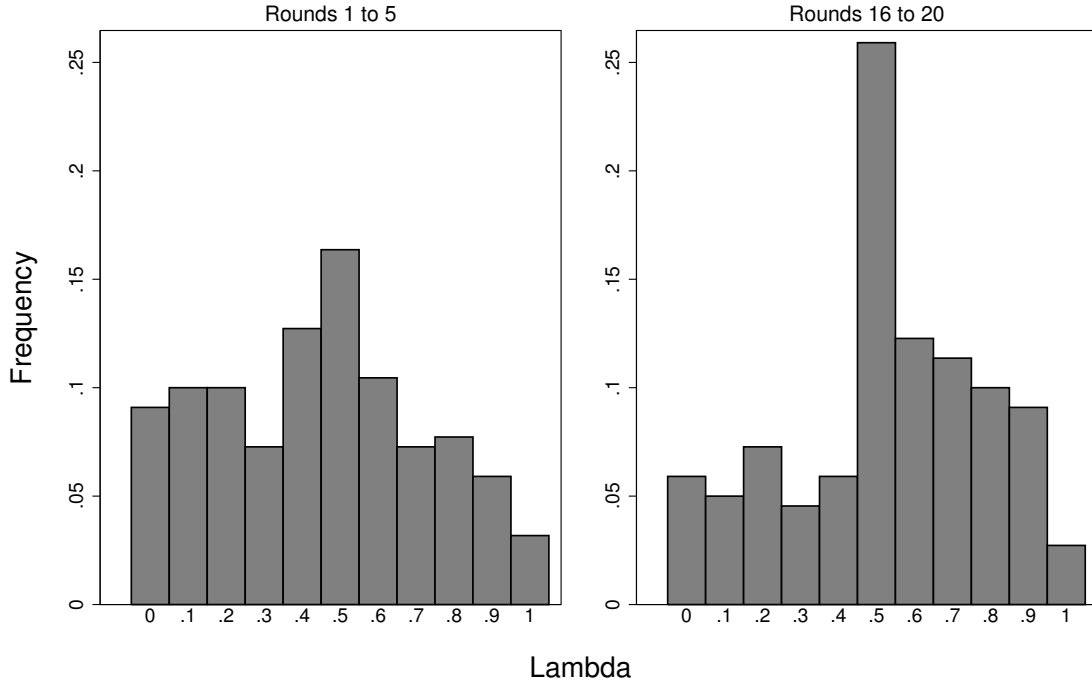
Figure A1: λ Choices by Round in the Short Sample

$D = 0.24, p < 0.01$ 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 ($D = 0.26, P < 0.01$).

C   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 λ. We estimate the causal impact
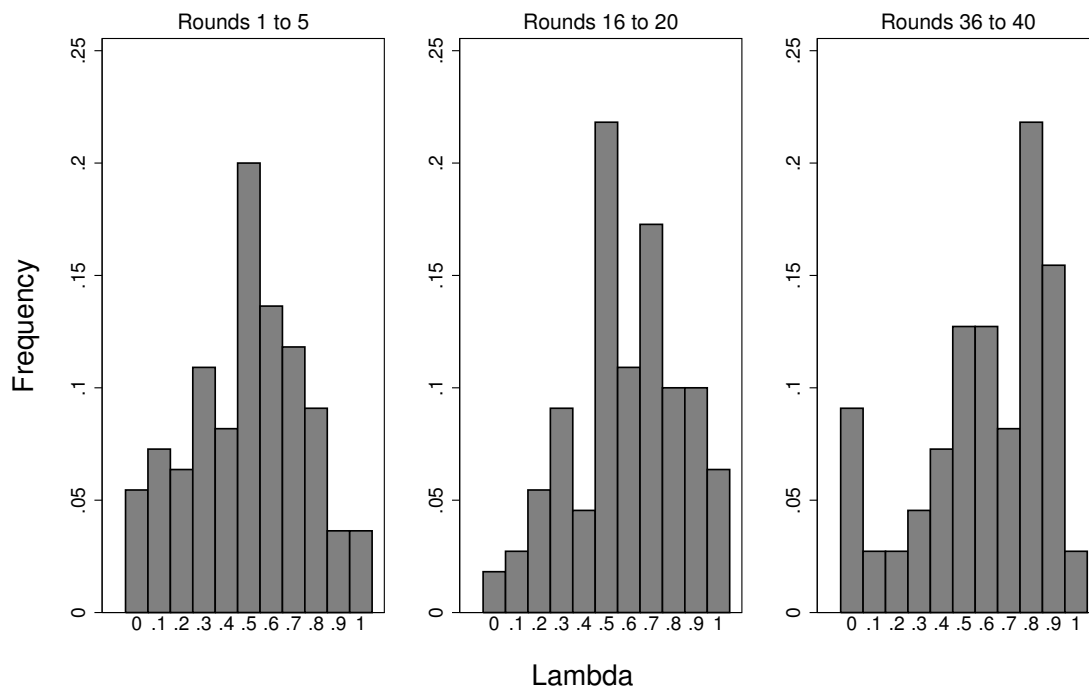
Figure A2: λ Choices by Round in the Long Sample

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 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.[18] 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 λ for this. Here λ can be thought of as a cost of cooperation, and random variation in λ can thus lead to random variation in cooperation.

The difference in how λ 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 λ as an instrument for whether an individual cooperates in both stages of a

---

[18]Another similar situation comes from Ham, Kagel and Lehrer (2005) in which the researchers study the impact of cash balances in auction behavior, using randomness in previous rounds as an instrument for cash balances.

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 λ that they did not choose. We then use λ 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 λ 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 λ was implemented in the current game.

Maximizing the relevance of our instruments requires a different functional form across conditions. Figure A3 shows the relationship between λ 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 λ and joint payoffs fits well. In the endogenous condition, an indicator variable for whether λ 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.
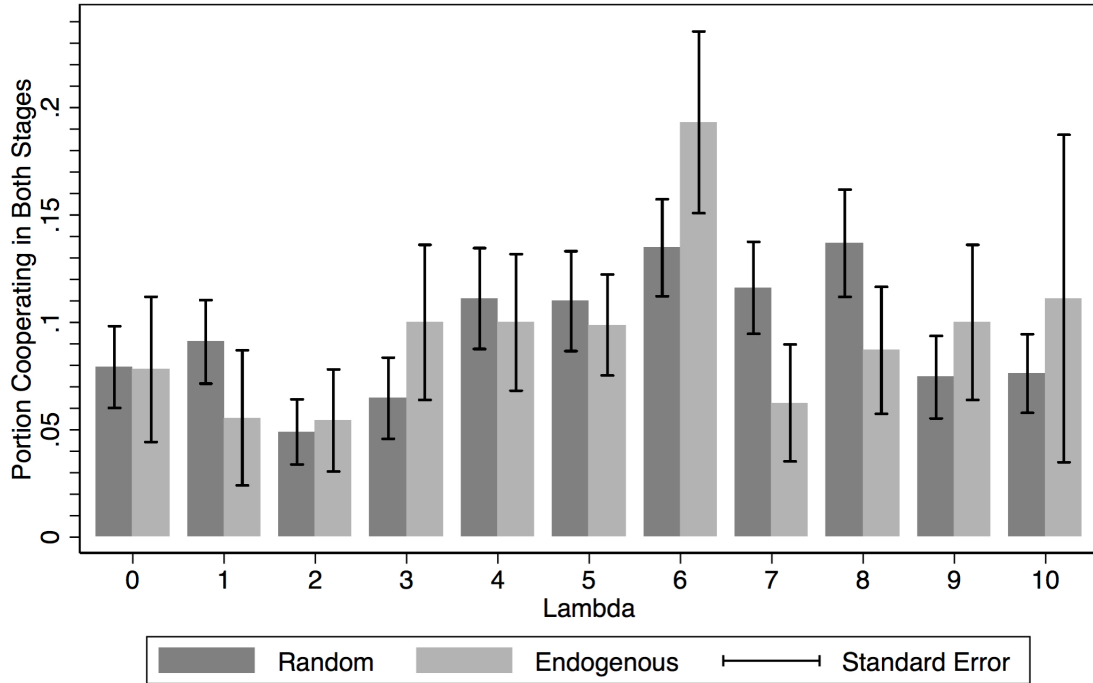

Figure A3: Cooperation and λ across Conditions

Our first-stage IV specifications are

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

in the random condition and

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

in the endogenous condition. $C^1_{i,t}$ and $C^2_{i,t}$ 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 λ_{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

$$\text{IV-R2:}\ 1(C^1_{i,t}, C^2_{i,t} = 1) = \alpha_i + \gamma_t + \zeta * 1(C^1_{i,t-1}, C^2_{i,t-1} = 1) + \eta\hat{\epsilon}_{i,j,t-1} +$$

$$\beta_1\lambda_t + \beta_2\lambda_t^2 + \beta_3\lambda_t^3 + \delta_1 C^1_{k,t} + \delta_2 C^2_{j,t-1} + \delta_3 C^1_{j,t-1} + \epsilon_{i,j,k,t}$$

In the endogenous condition, we introduce an additional control for whether individual i's choice of λ is implemented in round t, $L_{i,t}$. The sample is again restricted to those who did not choose λ in the previous game.

$$\text{IV-E2:}\ 1(C^1_{i,t}, C^2_{i,t} = 1) = \alpha_i + \gamma_t + \zeta * 1(C^1_{i,t-1}, C^2_{i,t-1} = 1) + \eta\hat{\epsilon}_{i,j,t-1} +$$

$$\beta * 1(\lambda_t = 0.6) + \delta_1 C^1_{k,t} + \delta_2 C^2_{j,t-1} + \delta_3 C^1_{j,t-1} + \theta L_{i,t} + \epsilon_{i,j,k,t}\ .$$

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.[19]

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 λ 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 $p < 0.01$.

---

[19]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.

Our estimates of reinforcement learning are found in Table A2. OLS estimates of the relationship between lagged cooperation and present cooperation yield similar results in both conditions, and the effect is positive and significant.[20] 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.

Table A2: IV Estimates of Reinforcement Learning in Cooperation

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

[20]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.

D   Subjects' Instructions: Random Condition

Welcome to the Economics Study

Welcome

Thank you for participating in this study. We expect this study to last about 90 minutes. Your earnings in this study will be paid to you in cash at the end of the session.

Throughout the experiment your identity will be kept totally private. Neither the experimenter nor the other participants will ever be able to tie you to your decisions.

The Experiment

In this experiment you will play a series of 20 **games**. In each of the 20 games you will be randomly paired with one other person for that game. Your partner in each game will change randomly throughout the study. You will never be able to predict which of the other participants in the room you are paired with for any game. Also, you will *never* play anyone more than one time.

In each game, you and your partner will make choices in two **rounds**. When the two rounds are over, your game will be complete. Then you will be randomly assigned a new partner and start a new game, again with two rounds.

You will repeat this process until you have completed a total of 20 games. Since each game will have two rounds, you will be in a total of 40 rounds over the course of the experiment.

In each game you will earn **tokens**. The tokens you earn in each game will be deposited in your Earnings Account. At the end of the study you will be paid $0.06 for every chip in your Earnings Account.

Each Round

Each game has two rounds. In each round you will decide between one of two options. You can either **pull** an amount X to yourself, or you can **push** an amount Y to your partner. In every decision, the amount you can push is three times the amount you can pull, that is, Y=3X. However, the values of X and Y will be changing from round to round.

Here is an example of a decision:

I choose to:	○ **pull 10 tokens** to myself, or	My partner chooses to:	○ pull 10 tokens to him/herself, or
	○ **push 30 tokens** to the other player		○ push 30 tokens to the other player
Submit			

There are four possible outcomes:

Possible Outcome 1: If you decide to pull 10 tokens to yourself and your partner decides to push 30 tokens to

you, then your payoff is 40 tokens and your partner's payoff is 0 tokens.

Possible Outcome 2: If you decide to pull 10 tokens to yourself but instead your partner decides to pull 10 tokens for himself, then your payoff is 10 tokens and your partner's payoff is 10 tokens.

Possible Outcome 3: If you decide to push 30 tokens to your partner and your partner decides to push 30 tokens to you, then your payoff is 30 tokens and your partner's payoff is 30 tokens.

Possible Outcome 4: If you decide to push 30 tokens to your partner but instead your partner decides to pull 10 tokens to himself, then your payoff is 0 tokens and your partner's payoff is 40 tokens.

As you can see, your partner will be faced with the same decision as you. You will both make your decisions at the same time. That is, you must make your decision without knowing what your partner is deciding.

Each Game

Each time you are paired with a new partner you will play a 2-round game with that person. In each round you will make a decision like that above.

Here is an example of what a game could look like:

Round 1 Decision:

Round 1 - Make a Choice	
I choose to: ○ **pull 3 tokens** to myself, or ○ **push 9 tokens** to the other player Submit	My partner chooses to: ○ pull 3 tokens to him/herself, or ○ push 9 tokens to the other player
Round 2 - Next Round	
I choose to: ○ pull 7 tokens to myself, or ○ push 21 tokens to the other player	My partner chooses to: ○ pull 7 tokens to him/herself, or ○ push 21 tokens to the other player

Notice that when you are asked to make your decision in the first round, you will also be able to see the decision to be made in the second round. This is shown in the grayed-out portion of the decision screen.

So, for example, suppose that in Round 1, you decide to push 9 tokens to your partner and your partner also decides to push 9 tokens to you. Then your payoff for the round would be 9 tokens and your partner's payoff would also be 9 tokens.

You will be able to see the results of your decision and your partner's decision before you make your decision for the second round. The screen you will see for your second-round decision looks like this:

Round 1 - Results	
I choose to: ○ **pull 3 tokens** to myself, or ⦿ **push 9 tokens** to the other player	My partner chooses to: ○ pull 3 tokens to him/herself, or ⦿ push 9 tokens to the other player
Round 1 - Earnings	
You: 0 + 9 = 9	Your Partner: 0 + 9 = 9
Round 2 - Make A Choice	
I choose to: ○ **pull 7 tokens** to myself, or ○ **push 21 tokens** to the other player	My partner chooses to: ○ pull 7 tokens to him/herself, or ○ push 21 tokens to the other player

Submit

After seeing these results, you can go on to make a choice for Round 2. Suppose in this Round 2 you chose to push 21 tokens while your partner chose to pull 7. Then for this decision you will earning nothing while your partner earns 7 + 21 = 28 tokens.

This makes your total earnings for the game 9 + 0 = 9, while your partner's total earnings are 9 + 28 = 37. The results of this game will be reported to you like this:

Round 1 - Results	
I choose to: ○ **pull 3 tokens** to myself, or ⦿ **push 9 tokens** to the other player	My partner chooses to: ○ pull 3 tokens to him/herself, or ⦿ push 9 tokens to the other player
Round 1 - Earnings	
You: 0 + 9 = 9	Your Partner: 0 + 9 = 9
Round 2 - Make A Choice	
I choose to: ○ pull 7 tokens to myself, or ⦿ push 21 tokens to the other player	My partner chooses to: ⦿ pull 7 tokens to him/herself, or ○ push 21 tokens to the other player
Round 2 - Earnings	
You: 0 + 0 = 0	Your Partner: 7 + 21 = 28

That was the end of game 1.

When you finish viewing the results of the game, you can click ⌈ Next Game ⌋. Then you will be randomly assigned a new partner from the others in the room and begin a new 2-round game.

How the Amounts to Pull and Push will Change

The amounts available to pull and push will change from round-to-round and from game-to-game. Here we explain how these values will be set.

For each decision, the number of tokens available to push will always be 3 times the number available to pull. For example, if you can pull 2 then you can push 6. Or, if you can pull 8, then you can push 24. If you can pull 10 then you can push 30.

In any decision the tokens you can pull will always be between 0 and 10. Since the push amounts are three times the pull amounts, the amount you can push will always be between 0 and 30.

There will also be a special way the pull and push amounts are determined within a game. In particular, the number of tokens you can pull in Round 1 plus the number you can pull in Round 2 will always equal 10. For example, if you can pull 4 in Round 1 then you can pull 6 in Round 2. Or, if you can pull 1 in Round 1 then you can pull 9 in Round 2. If you can pull 10 in Round 1, then you can pull 0 in Round 2.

Note that since the pull amounts in Round 1 and Round 2 always sum to 10, this means that the push amounts in the two rounds will always sum to 30. In other words, all games will have the same feature that the total amount to pull across the the two rounds is 10 and the total amount to push is 30. How games will differ is in how many push and pull tokens are allocated to Round 1 and how many to Round 2.

Finally, the push and pull amounts you see in any game will be drawn at random from all the possible pull and push amounts that meet these rules above. You will never know what values you will see in future games, but all possible values are equally likely.

So there are three things to remember about how the pull and push amounts are set:

1. The push amounts are always 3 times the pull amounts.
2. In each game the pull amount in Round 1 and the pull amount in Round 2 always sume to 10. As a result, the push amount in Round 1 and the push amount in Round 2 sum to 30.
3. The values in each game are determined at random from all the values that meet rules (1) and (2)

Your History

If you want to look back at the history of play you have seen over the experiment, you can do this from any screen by hitting the button ⌈ View My History ⌋. This will show you your decisions, your partner's decision, and your earnings in each game.

Overview of the Experiment

As we are about to begin, keep these things in mind:

- You will play a total of 20 2-round games.
- For each 2-round game, you will play with the same partner for both of the rounds.
- When you start a new game, you will get a new partner, chosen at random from everyone here today.
- You will never play the same person more than once.
- In each 2-round game the total amount to pull across the two rounds is 10 and the total amount to push is 30. The games will differ in how much of this is allocated to Round 1 and how much to Round 2.
- You will be paid your total earnings across all 20 of the 2-round games.
- Each token you earn is worth $0.06.
- The experiment will last about 90 minutes.

Thanks for participating. Good luck!

Begin!

E   Subjects' Instructions: Endogenous Condition

Welcome to the Economics Study

Welcome

Thank you for participating in this study. We expect this study to last about 90 minutes. Your earnings in this study will be paid to you in cash at the end of the session.

Throughout the experiment your identity will be kept totally private. Neither the experimenter nor the other participants will ever be able to tie you to your decisions.

The Experiment

In this experiment you will play a series of 3 **games**. In each of the 3 games you will be randomly paired with one other person for that game. Your partner in each game will change randomly throughout the study. You will never be able to predict which of the other participants in the room you are paired with for any game. Also, you will *never* play anyone more than one time.

In each game, you and your partner will make choices in two **rounds**. When the two rounds are over, your game will be complete. Then you will be randomly assigned a new partner and start a new game, again with two rounds.

You will repeat this process until you have completed a total of 3 games. Since each game will have two rounds, you will be in a total of 6 rounds over the course of the experiment.

In each game you will earn **tokens**. The tokens you earn in each game will be deposited in your Earnings Account. At the end of the study you will be paid $0.04 for every chip in your Earnings Account.

Each Round

Each game has two rounds. In each round you will decide between one of two options. You can either **pull** an amount X to yourself, or you can **push** an amount Y to your partner. In every decision, the amount you can push is three times the amount you can pull, that is, Y=3X. However, the values of X and Y will be changing from round to round.

Here is an example of a decision:

I choose to: ○ **pull 10 tokens** to myself, or ○ **push 30 tokens** to the other player ⬚ Submit	My partner chooses to: ○ pull 10 tokens to him/herself, or ○ push 30 tokens to the other player

There are four possible outcomes:

Possible Outcome 1: If you decide to pull 10 tokens to yourself and your partner decides to push 30 tokens to you, then your payoff is 40 tokens and your partner's payoff is 0 tokens.

Possible Outcome 2: If you decide to pull 10 tokens to yourself but instead your partner decides to pull 10 tokens for himself, then your payoff is 10 tokens and your partner's payoff is 10 tokens.

Possible Outcome 3: If you decide to push 30 tokens to your partner and your partner decides to push 30 tokens to you, then your payoff is 30 tokens and your partner's payoff is 30 tokens.

Possible Outcome 4: If you decide to push 30 tokens to your partner but instead your partner decides to pull 10 tokens to himself, then your payoff is 0 tokens and your partner's payoff is 40 tokens.

As you can see, your partner will be faced with the same decision as you. You will both make your decisions at the same time. That is, you must make your decision without knowing what your partner is deciding.

Each Game

Each time you are paired with a new partner you will play a 2-round game with that person. In each round you will make a decision like that above.

Here is an example of what a game could look like:

Round 1 Decision:

Round 1 - Make a Choice	
I choose to: ○ **pull 3 tokens** to myself, or ○ **push 9 tokens** to the other player 〔Submit〕	My partner chooses to: ○ pull 3 tokens to him/herself, or ○ push 9 tokens to the other player
Round 2 - Next Round	
I choose to: ○ pull 7 tokens to myself, or ○ push 21 tokens to the other player	My partner chooses to: ○ pull 7 tokens to him/herself, or ○ push 21 tokens to the other player

Notice that when you are asked to make your decision in the first round, you will also be able to see the decision to be made in the second round. This is shown in the grayed-out portion of the decision screen.

So, for example, suppose that in Round 1, you decide to push 9 tokens to your partner and your partner also decides to push 9 tokens to you. Then your payoff for the round would be 9 tokens and your partner's payoff would also be 9 tokens.

You will be able to see the results of your decision and your partner's decision before you make your decision for the second round. The screen you will see for your second-round decision looks like this:

Round 1 - Results	
I choose to: ○ **pull 3 tokens** to myself, or ⊙ **push 9 tokens** to the other player	My partner chooses to: ○ pull 3 tokens to him/herself, or ⊙ push 9 tokens to the other player
Round 1 - Earnings	
You: 0 + 9 = 9	Your Partner: 0 + 9 = 9
Round 2 - Make A Choice	
I choose to: ○ **pull 7 tokens** to myself, or ○ **push 21 tokens** to the other player 〔Submit〕	My partner chooses to: ○ pull 7 tokens to him/herself, or ○ push 21 tokens to the other player

After seeing these results, you can go on to make a choice for Round 2. Suppose in this Round 2 you chose to push 21 tokens while your partner chose to pull 7. Then for this decision you will earning nothing while your partner earns 7 + 21 = 28 tokens.

This makes your total earnings for the game 9 + 0 = 9, while your partner's total earnings are 9 + 28 = 37. The results of this game will be reported to you like this:

Round 1 - Results	
I choose to: ○ **pull 3 tokens** to myself, or ⊙ **push 9 tokens** to the other player	My partner chooses to: ○ pull 3 tokens to him/herself, or ⊙ push 9 tokens to the other player

Round 1 - Earnings	
You: 0 + 9 = 9	Your Partner: 0 + 9 = 9

Round 2 - Make A Choice	
I choose to: ○ pull 7 tokens to myself, or ⊙ push 21 tokens to the other player	My partner chooses to: ⊙ pull 7 tokens to him/herself, or ○ push 21 tokens to the other player

Round 2 - Earnings	
You: 0 + 0 = 0	Your Partner: 7 + 21 = 28

That was the end of game 1.

Total Game Earnings
You: 9 tokens   Your Partner: 37 tokens

When you finish viewing the results of the game, you can click [Next Game]. Then you will be randomly assigned a new partner from the others in the room and begin a new 2-round game.

How the Amounts to Pull and Push will Change

The amounts available to pull and push will change from round-to-round and from game-to-game. Here we explain how these values will be set.

For each decision, the number of tokens available to push will always be 3 times the number available to pull. For example, if you can pull 2 then you can push 6. Or, if you can pull 8, then you can push 24. If you can pull 10 then you can push 30.

In any decision the tokens you can pull will always be between 0 and 10. Since the push amounts are three times the pull amounts, the amount you can push will always be between 0 and 30.

There will also be a special way the pull and push amounts are determined within a game. In particular, the number of tokens you can pull in Round 1 plus the number you can pull in Round 2 will always equal 10. For example, if you can pull 4 in Round 1 then you can pull 6 in Round 2. Or, if you can pull 1 in Round 1 then you can pull 9 in Round 2. If you can pull 10 in Round 1, then you can pull 0 in Round 2.

Note that since the pull amounts in Round 1 and Round 2 always sum to 10, this means that the push amounts in the two rounds will always sum to 30. In other words, all games will have the same feature that the total amount to pull across the the two rounds is 10 and the total amount to push is 30. How games will differ is in how many push and pull tokens are allocated to Round 1 and how many to Round 2.

In each round the push and pull values will be set by _one of_ the two players. Before any round, all players will choose which of the 11 possible push and pull values they would like to play. Then after you are paired with another player, the computer will randomly select either the push and pull values that you chose, or the push and pull values that the other player chose.

You will choose the push and pull values that you wish to play by filling out a form like this below. Try it to see how it works.

Preliminary Round: Select the game to be played next round. Your partner will also be selecting a game to play. Which game you actually play will be determined at random to be either the game you chose or the game your partner chose. You and your partner must always play the same game.

Select a number below in order to set the pull and push values in the game you will play next:

○0 ○1 ○2 ◉3 ○4 ○5 ○6 ○7 ○8 ○9 ○10

Submit your decision when you have selected the game below that you wish to play next: [Submit]

Round 1

I choose to: ○ pull 3 tokens to myself, or ○ push 9 tokens to the other player	My partner chooses to: ○ pull 3 tokens to him/herself, or ○ push 9 tokens to the other player

Round 2

I choose to: ○ pull 7 tokens to myself, or ○ push 21 tokens to the other player	My partner chooses to: ○ pull 7 tokens to him/herself, or ○ push 21 tokens to the other player

---

Results from the Preliminary Stage:

You Chose: ○0 ○1 ○2 ◉3 ○4 ○5 ○6 ○7 ○8 ○9 ○10

The computer chose randomly between your choice and the choice of your partner. The result is that both you and your partner will play this game:

○0 ○1 ○2 ○3 ○4 ○5 ◉6 ○7 ○8 ○9 ○10

Begin Round 1:

Round 1 - Make a Choice

I choose to: ◉ **pull 6 tokens** to myself, or ○ **push 18 tokens** to the other player [Submit]	My partner chooses to: ○ pull 6 tokens to him/herself, or ○ push 18 tokens to the other player

Round 2 - Next Round

I choose to: ○ pull 4 tokens to myself, or ○ push 12 tokens to the other player	My partner chooses to: ○ pull 4 tokens to him/herself, or ○ push 12 tokens to the other player

---

Results from the Preliminary Stage:

You Chose: ○0 ○1 ○2 ◉3 ○4 ○5 ○6 ○7 ○8 ○9 ○10

The computer chose randomly between your choice and the choice of your partner. The result is that both you and your partner will play this game:

○0 ○1 ○2 ○3 ○4 ○5 ◉6 ○7 ○8 ○9 ○10

Begin Round 2:

Round 1 - Results	
I choose to: ○ **pull 6 tokens** to myself, or ◉ **push 18 tokens** to the other player	My partner chooses to: ○ pull 6 tokens to him/herself, or ◉ push 18 tokens to the other player
Round 1 - Earnings	
You: 0 + 18 = 18	Your Partner: 0 + 18 = 18
Round 2 - Make A Choice	
I choose to: ◉ **pull 4 tokens** to myself, or ○ **push 12 tokens** to the other player	My partner chooses to: ○ pull 4 tokens to him/herself, or ○ push 12 tokens to the other player
[Submit]	

Results from the Preliminary Stage:

You Chose: ○0 ○1 ○2 ◉3 ○4 ○5 ○6 ○7 ○8 ○9 ○10

The computer chose randomly between your choice and the choice of your partner. The result is that both you and your partner will play this game:

○0 ○1 ○2 ○3 ○4 ○5 ◉6 ○7 ○8 ○9 ○10

Round 1 - Results	
I choose to: ○ **pull 6 tokens** to myself, or ◉ **push 18 tokens** to the other player	My partner chooses to: ○ pull 6 tokens to him/herself, or ◉ push 18 tokens to the other player
Round 1 - Earnings	
You: 0 + 18 = 18	Your Partner: 0 + 18 = 18
Round 2 - Make A Choice	
I choose to: ○ pull 4 tokens to myself, or ◉ push 12 tokens to the other player	My partner chooses to: ◉ pull 4 tokens to him/herself, or ○ push 12 tokens to the other player
Round 2 - Earnings	
You: 0 + 0 = 0	Your Partner: 4 + 12 = 16

That was the end of game 1.

Total Game Earnings

You: 18 tokens   Your Partner: 34 tokens

So there are three things to remember about how the pull and push amounts are set:

1. The push amounts are always 3 times the pull amounts.
2. In each game the pull amount in Round 1 and the pull amount in Round 2 always sume to 10. As a result, the push amount in Round 1 and the push amount in Round 2 sum to 30.
3. Before any game, both players will play a Preliminary round where they choose the push and pull values for the game they wish to play. Which push and pull values you actually play will be determined at random to be either the those you chose or the those your partner chose.

Your History

If you want to look back at the history of play you have seen over the experiment, you can do this from any screen by hitting the button [View My History]. This will show you your decisions, your partner's decision, and your earnings in each game.

Overview of the Experiment

As we are about to begin, keep these things in mind:

- You will play a total of 20 2-round games.
- For each 2-round game, you will play with the same partner for both of the rounds.
- When you start a new game, you will get a new partner, chosen at random from everyone here today.
- You will never play the same person more than once.
- In each 2-round game the total amount to pull across the two rounds is 10 and the total amount to push is 30. The games will differ in how much of this is allocated to Round 1 and how much to Round 2.
- You will be paid your total earnings across all 20 of the 2-round games.
- Each token you earn is worth $0.06.
- The experiment will last about 90 minutes.

Thanks for participating. Good luck!