

ORIGINAL ARTICLE

Building rational cooperation on their own: Learning to start small

James Andreoni¹ | Michael A. Kuhn²  | Larry Samuelson³ 

¹Department of Economics, University of California, San Diego, La Jolla, California & NBER

²Department of Economics, University of Oregon, Eugene, Oregon

³Department of Economics, Yale University, New Haven, Connecticut

Correspondence

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

Funding information

The National Science Foundation (SES-1427355);
UC San Diego Economics Department;
The National Science Foundation (SES-1459158).

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.

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

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

increasing stakes plays a key role in supporting cooperation. Rauch and Watson (2003) present empirical evidence that starting small plays a role in building 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 Kamiyo, Ozono, and Shimizu (2016) and Ozono, Kamiyo, 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, whereas 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’ 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.

² Others who 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 Ehrhart and Keser (1999), Cinyabuguma, Page, and Putterman (2005), Page, Putterman, and Unel (2005), and Ahn, Isaac, and Salmon (2008, 2009).

		Stage 1		Stage 2	
		C	D	C	D
C	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, $0 \leq x_1 \leq 10$ and $x_1 + x_2 = 10$

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 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 five sessions each involving 22 subjects playing 20 rounds of the twice-repeated prisoners' dilemma, implying 1,100 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 & 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.

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 toward 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.

⁵ All data were 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 Supporting Information.

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.




FIGURE 2 Payoffs by value of λ and condition

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 .

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 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 (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

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 rescaling 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 (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 the online Supporting Information Section A, Table A1.

TABLE 1 Relationship between λ and payoffs across conditions

	Sample restriction		
	All rounds	Rounds 1–10	Rounds 11–20
<i>Random condition terms:</i>			
λ	-1.408 (6.693)	-12.623 (11.466)	-3.205 (6.949)
λ^2	43.823*** (15.883)	66.310** (26.457)	49.563*** (15.088)
λ^3	-44.768*** (10.687)	-59.212*** (17.495)	-46.765*** (9.463)
<i>Endogenous interactions:</i>			
$\lambda * 1(\text{endogenous} = 1)$	4.186 (12.847)	19.080 (24.381)	5.619 (16.906)
$\lambda^2 * 1(\text{endogenous} = 1)$	-6.602 (31.222)	-57.092 (61.421)	2.515 (38.373)
$\lambda^3 * 1(\text{endogenous} = 1)$	4.819 (21.943)	46.462 (42.312)	-10.990 (25.209)
H_0 : Endog. interactions = 0	$F(3, 1361) = 0.28$ ($p = 0.84$)	$F(3, 601) = 1.07$ ($p = 0.36$)	$F(3, 601) = 1.09$ ($p = 0.35$)
Random payoff-max. λ_r^*	0.636 (0.020)	0.635 (0.029)	0.673 (0.019)
Endogenous payoff-max. λ_e^*	0.656 (0.030)	0.717 (0.178)	0.623 (0.031)
H_0 : $\lambda_r^* = \lambda_e^*$	$\chi^2(1) = 0.31$ ($p = 0.58$)	$\chi^2(1) = 0.21$ ($p = 0.65$)	$\chi^2(1) = 1.76$ ($p = 0.19$)
N	1,540	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 nonlinear 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.

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 Subsection 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 the online Supporting Information Section B, along with statistical tests of the difference in distributions.

We partition our sample into three intervals. Call rounds 1–6 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, whereas 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

TABLE 2 Evolution of subsample sizes over time

Group, Rounds	Fraction of sample in each subsample		
	Start small	Even	Start large
<i>Panel A: Short sample</i>			
<i>Unconditional groups:</i>			
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
<i>Conditional on starting small in beginning:</i>			
End (Rounds 15–20 Rounds 1–6 = SS)	0.63	0.19	0.19
<i>Conditional on starting large in beginning:</i>			
End (Rounds 15–20 Rounds 1–6 = SL)	0.52	0.04	0.44
N = 44 per round			
<i>Panel B: Long sample</i>			
<i>Unconditional groups:</i>			
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
<i>Conditional on starting small in beginning:</i>			
End (Rounds 35–40 Rounds 1–6 = SS)	0.70	0.00	0.30
<i>Conditional on starting large in beginning:</i>			
End (Rounds 35–40 Rounds 1–6 = SL)	0.78	0.00	0.22
N = 22 per round			

groups gravitate to starting small by the end, and those who started small at the beginning do so to an even greater degree.¹¹

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 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 toward 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 the online Supporting Information 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 side 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.

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.

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 ($p = 0.10$).

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 $	$\Pr(\lambda_{i,t} \in [0.6, 0.9])$	<i>H</i> -index
<i>Panel A: Short sample</i>					
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
<i>Panel B: Long sample</i>					
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 overestimating the degree of choice dispersion.

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 analog to Table 2, we use an indicator variable for starting small as an outcome variable. Last, 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 four of five 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

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 λ 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.

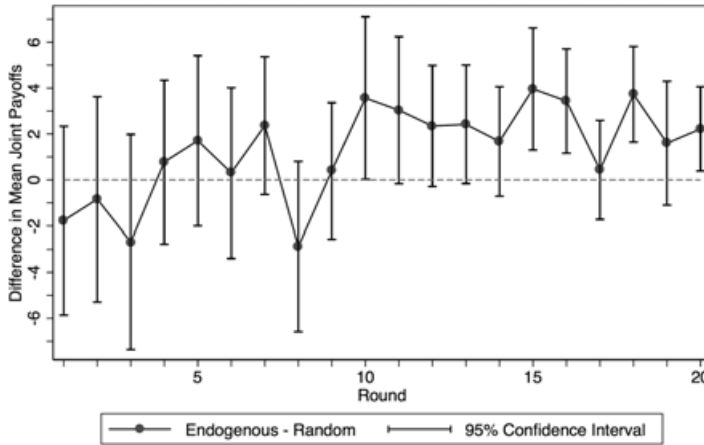


FIGURE 3 Excess payoffs in endogenous condition over random condition by round

TABLE 4 Payoff differences across conditions

	Sample restriction					
	All rounds			Rounds 11–20		
	(1)	(2)	(3)	(4)	(5)	(6)
λ (<i>endogenous</i> = 1)	1.290 (1.190)	1.952 (1.679)	-0.233 (0.947)	2.485** (1.142)	4.396** (1.990)	-0.400 (0.790)
Constant	28.374 (0.553)	29.335 (1.130)	25.620 (1.171)	24.442 (0.528)	23.663 (1.120)	21.073 (0.715)
λ fixed effects?	No	Yes	Yes	No	Yes	Yes
λ range	All	Central	Extreme	All	Central	Extreme
N	1,540	446	1,074	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.

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 show 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.

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: (a) the causal effect of a value of λ being endogenously rather than randomly derived, and (b) 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)

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

TABLE 5 Effect of λ on cooperation in the endogenous condition

	First-stage cooperation	
	All rounds	Rounds 11–20
<i>Central/Used</i>	0.024	0.112
	(0.079)	(0.102)
<i>Central/Unused</i>	0.021	0.333**
	(0.090)	(0.124)
Difference: <i>Used–Unused</i>	0.002	–0.221*
	(0.068)	(0.125)
<i>Extreme/Used</i>	–0.012	0.012
	(0.043)	(0.070)
<i>Extreme/Unused</i>	–0.122**	–0.103
	(0.050)	(0.076)
Difference: <i>Used–Unused</i>	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.

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 could be warning signs, and we might expect to see a negative effect of the endogenous condition, whereas, 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 noncentral values), a subject should be *less* likely to cooperate when their choice is *not* selected for 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: (a) *Central/Used* = 1 if an individual's central choice of λ was implemented; (b) *Central/Unused* = 1 if an individual's central choice of λ was not implemented and a different central λ was implemented; (c) *Extreme/Used* = 1 when an individual's extreme choice of λ was implemented; and (d) *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.

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, whereas 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 λ 's 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, whereas 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

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 the 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. In addition, 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 prosocially in social dilemmas, largely based on moral principles or altruistic 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.

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

ACKNOWLEDGMENT

We appreciate advice from Doug Miller and David Eil.

ORCID

Michael A. Kuhn  <http://orcid.org/0000-0001-9850-1273>

Larry Samuelson  <http://orcid.org/0000-0002-2702-2555>

REFERENCES

- Ahn, T. K., Isaac, R. M., & Salmon, T. C. (2008). Endogenous group formation. *Journal of Public Economic Theory*, 10, 171–194.
- Ahn, T. K., Isaac, R. M., & Salmon, T. C. (2009). Coming and going: Experiments on endogenous group sizes for excludable public goods. *Journal of Public Economics*, 93, 336–351.
- Ali, S. N., & Miller, D. A. (2013). *Enforcing cooperation in networked societies*. Working paper.
- Andreoni, J. (2017). *Satisfaction guaranteed: When moral preferences meet moral hazard*. Working paper.
- Andreoni, J., & Gee, L. (2012). Gun for hire: Delegated enforcement and peer punishment in public goods provision. *Journal of Public Economics*, 96, 1036–1046.
- Andreoni, J., & Miller, J. (2002). Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70, 737–753.
- Andreoni, J., & Samuelson, L. (2006). Building rational cooperation. *Journal of Economic Theory*, 127, 117–154.
- Andreoni, J., & Varian, H. (1999). Preplay contracting in the prisoners dilemma. *Proceedings of the National Academy of Sciences*, 96, 10933–10938.
- Binmore, K., Proulx, C., Samuelson, L., & Swierzbinski, J. (1998). Hard bargains and lost opportunities. *Economic Journal*, 108, 1279–1298.
- Blonski, M., & Probst, D. (2004). *The emergence of trust*. Working paper.
- Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton, NJ: Russell Sage Foundation and Princeton University Press.
- Charness, G., & Yang, C.-L. (2014). Starting small toward voluntary formation of efficient large groups in public goods provision. *Journal of Economic Behavior and Organization*, 102, 119–132.
- Cinyabuguma, M., Page, T., & Putterman, L. (2005). Cooperation under the threat of expulsion in a public goods experiment. *Journal of Public Economics*, 89, 1421–1435.
- Datta, S. (1996). *Building trust* (STICERD Discussion Paper 305).
- Diamond, D. W. (1989). Reputation acquisition in debt markets. *Journal of Political Economy*, 97, 828–862.
- Duffy, J., Ochs, J., & Vesterlund, L. (2007). Giving little by little: Dynamic voluntary contribution games. *Journal of Public Economics*, 91, 1708–1730.
- Ehrhart, K.-M., & Keser, C. (1999). *Mobility and cooperation: On the run*. Working paper.
- Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, 90, 980–994.
- Fujiwara, T., Meng, K. C., & Vogl, T. (2013). *Estimating habit formation in voting*. NBER Working Paper No. 19721.
- Ghosh, P., & Ray, D. (1996). Cooperation in community interaction without information flows. *Review of Economic Studies*, 63, 491–519.
- Greif, A., & Laitin, D. D. (2004). A theory of endogenous institutional change. *American Political Science Review*, 98, 633–652.
- Ham, J. C., Kagel, J. H., & Lehrer, S. F. (2005). Randomization, endogeneity and laboratory experiments: The role of cash balances in private value auctions. *Journal of Econometrics*, 125, 175–205.
- Kamijo, Y., Ozono, H., & Shimizu, K. (2016). Overcoming coordination failure using a mechanism based on gradualism and endogeneity. *Experimental Economics*, 19(1), 202–217.
- Kocher, M., Martinsson, P., & Visser, M. (2012). Social background, cooperative behavior and norm enforcement. *Journal of Economic Behavior and Organization*, 81, 341–354.
- Kranton, R. E. (1996). The formation of cooperative relationships. *Journal of Law, Economics and Organization*, 12(1), 214–233.
- Kreps, D. M., Milgrom, P., Roberts, J., & Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, 27, 245–252.

- Mailath, G. J., & Samuelson, L. (2006). *Repeated games and reputations: Long-run relationships*. Oxford: Oxford University Press.
- Marx, L. M., & Matthews, S. A. (2000). Dynamic voluntary contribution to a public project. *Review of Economic Studies*, 67, 327–358.
- Nikiforakis, N. (2008). Punishment and counter-punishment in public good games: Can we really govern ourselves? *Journal of Public Economics*, 92, 91–112.
- Ozono, H., Kamijo, Y., & Shimizu, K. (2016). Institutionalize reciprocity to overcome the public goods provision problem. *PLoS ONE*, 11(6), e0154321.
- Page, T., Putterman, L., & Unel, B. (2005). Voluntary association in public goods experiments: Reciprocity, mimicry and efficiency. *Economic Journal*, 115(506), 1032–1053.
- Peters, M. (1995). Decentralized markets and endogenous institutions. *Canadian Journal of Economics*, 28, 227–260.
- Rand, D. G., Dreber, A., Ellingsen, T., Fudenberg, D., & Nowak, M. A. (2009). Positive interactions promote public cooperation. *Science*, 325(5945), 1272–1275.
- Rauch, J. E., & Watson, J. (2003). Starting small in an unfamiliar environment. *International Journal of Industrial Organization*, 21, 1021–1042.
- Schelling, T. C. (1960). *Strategy of conflict*. Cambridge, MA: Harvard University Press.
- Sobel, J. (1985). A theory of credibility. *Review of Economic Studies*, 52, 557–573.
- Watson, J. (1999). Starting small and renegotiation. *Journal of Economic Theory*, 85, 52–90.
- Watson, J. (2002). Starting small and commitment. *Games and Economic Behavior*, 38(1), 176–199.
- Weber, R. A. (2006). Managing growth to achieve efficient coordination in large groups. *American Economic Review*, 96, 114–126.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

How to cite this article: Andreoni J, Kuhn MA, Samuelson L. Building rational cooperation on their own: Learn to start small. *J Public Econ Theory*. 2019;21:812–825. <https://doi.org/10.1111/jpet.12322>