

Long-Run Learning in Games of Cooperation

WINTER MASON, Facebook
SIDDHARTH SURI, Microsoft Research
DUNCAN J. WATTS, Microsoft Research

Cooperation in repeated games has been widely studied in experimental settings; however, the duration over which players participate in such experiments is typically confined to at most hours, and often to a single game. Given that in real world settings people may have years of experience, it is natural to ask how behavior in cooperative games evolves over the long run. Here we analyze behavioral data from three distinct games involving 571 individual experiments conducted over a two-year interval. First, in the case of a standard linear public goods game we show that as players gain experience, they become less generous both on average and in particular towards the end of each game. Second, we analyze a multiplayer prisoner's dilemma where players are also allowed to make and break ties with their neighbors, finding that experienced players show an increase in cooperativeness early on in the game, but exhibit sharper "endgame" effects. Third, and finally, we analyze a collaborative search game in which players can choose to act selfishly or cooperatively, finding again that experienced players exhibit more cooperative behavior as well as sharper endgame effects. Together these results show consistent evidence of long-run learning, but also highlight directions for future theoretical work that may account for the observed direction and magnitude of the effects.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics

Additional Key Words and Phrases: Cooperation; Social Networks; Learning; Prisoner's Dilemma; Public Goods

1. INTRODUCTION

The role of experience in social dilemmas has long been of interest to economists [Selten and Stoecker 1986; Andreoni 1988; Andreoni and Miller 1993; Ledyard 1995]. In brief, the interest stems from the difference between the predictions of standard economic theory and observed behavior in human subjects experiments. Theory predicts that in finitely repeated games of cooperation, rational players should defect on all turns via the familiar argument of backward induction [Osborne and Rubinstein 1994]. In contrast, experiments repeatedly show that a majority of human players tend to cooperate at first and, as the game is repeated, steadily decline their cooperation rates until the so-called end game approaches, yet cooperation levels rarely decline to zero [Ledyard 1995]. Kreps et al. [1982] proposed an ingenious solution to this puzzle—namely that players are rational, but have incomplete information about the rationality of others. In such a situation Kreps et al. [1982] showed that if players also believe with sufficiently large probability δ that other players employ a Tit-for-Tat strategy, and hence will cooperate at least until their partner defects, then it is in fact optimal for a rational player to cooperate at least until close to the end of the game. An interesting corollary of this result is in the presence of incomplete informa-

Author addresses: winteram@fb.com, suri@microsoft.com, duncan@microsoft.com

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'14, June 8–12, 2014, Stanford, CA, USA.

ACM 978-1-4503-2565-3/14/06 ...\$15.00.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

<http://dx.doi.org/10.1145/http://dx.doi.org/10.1145/2600057.2602892>

tion cooperation can dominate (for most of the game) even when all players are in fact rational.

The incomplete information hypothesis, however, raises an additional question: What happens to players' beliefs in the long run? Andreoni and Miller [1993] point out two possibilities. The first possibility is that if "true" altruism is in fact rare or absent—as Kreps et al. [1982] propose—then this fact will eventually become common knowledge, and rational players will respond by defecting earlier and earlier. In the long run, cooperation should disappear altogether just as predicted by the backward induction argument. The second possibility is that if Kreps et al. [1982] are wrong and genuine altruism really is present in sufficient degree, then as *this* knowledge becomes common, cooperation could be sustained for longer and longer as both rational cooperators and genuine altruists experience the benefit to deferring their eventual defection. Depending on the actual distribution of altruism in the population, in other words, the incomplete information hypothesis is consistent either with learning through experience leading to less cooperation or to more of it.

Alternative theoretical arguments, meanwhile, also make mixed predictions about the effect of experience on the tendency to cooperate. Andreoni [1995], for example, has argued separately that cooperation observed in experiments could arise because of genuinely altruistic beliefs that subjects bring with them from the outside world, but could also arise because subjects do not fully understand the incentives with which they are being presented. If the former applies, cooperation may persist indefinitely, whereas if the latter is the case, then subjects should learn over time not to cooperate. Although this explanation is theoretically distinct from the incomplete information hypothesis, it has the same consequence that, depending on the actual frequency of altruism in the population, cooperation could either increase or decrease over time as subjects gain experience.

Here we investigate long-run learning effects in three distinct games of cooperation, representing a total of 571 web-based experiments conducted over a period of two years and using a total of 466 unique individuals¹. As we will explain in more detail later (see Sections 3, 4, and 5), all three experiments were designed to answer research questions other than the one we address here: the first studied the effect of network topology on contribution in a linear public goods game [Suri and Watts 2011]; the second studied the effect of partner updating on cooperation in an iterated multi-player prisoner's dilemma game [Wang et al. 2012]; and the third studied the effect of network topology on collective learning among individuals searching a hidden "fitness landscape" where individuals were exposed to the previous locations and scores of their neighbors [Mason and Watts 2012]. To examine the effects of long-run learning, we exploit an artifact of the recruiting process used in all three experiments. That in order to attract and retain up to 24 individuals simultaneously for up to 30 minutes online², the researchers first recruited a panel of volunteers who were contacted repeatedly to participate in the experiments. An unintended consequence of this recruiting process was that a large number of subjects participated in multiple games and some participated in over 100. By examining how behavior varied with experience, both within and across individuals, we can therefore begin to address the question of long-run learning.

Although motivated by the incomplete information argument outline above, we emphasize that the experiments we study differ from the idealized construction of Kreps

¹Note: some players participated in more than one of the experiments we conducted, hence this total number is smaller than the sum of the three experiment populations, as given in later sections.

²All experiments were conducted using Amazon's Mechanical Turk, a crowdsourcing website that is increasingly popular with behavioral scientists for conducting human subjects experiments [Horton et al. 2011; Suri and Watts 2011; Paolacci et al. 2010; Mason and Suri 2012; Mason and Watts 2012; Wang et al. 2012].

et al. [1982] in at least two important ways. First, implicit in the Kreps et al. [1982] argument is the assumption that all players are present for the entire length of the “game,” and hence at each stage have the same experience³. As noted above, however, our experiments were characterized by considerable heterogeneity of experience: whereas a handful participated in over 100 experiments, the majority participated in only one. In such an environment, the state of incomplete information postulated by Kreps et al. [1982] might be prolonged indefinitely, leading either to slower rates of learning or possibly even preventing such learning from occurring at all. Further complicating matters, observed correlations between experience and cooperation might arise not because of learning at all, but simply on account of selection effects. That is individuals who play more games might be more or less predisposed towards cooperation than individuals who play few games, hence we must be careful to separate out within-individual effects from selection effects.

Second, while the first game that we consider (the repeated public goods games) bears a reasonably close resemblance to the repeated prisoner’s dilemma of Kreps et al. [1982], the other two games differ in potentially important ways. In particular, the Social Networking Game differs in that players were allowed to make and break ties, a mechanism that generated large and significant increases in cooperation levels [Wang et al. 2012]. The collaborative search game, meanwhile, although exhibiting the essential features of a social dilemma, is not framed explicitly as a game of cooperation, and also confronts the player with a rather more complicated set of choices than simply copy or explore. Precisely how important these differences are, and how they affect learning in the long run, is not at all clear from the theory. Investigating the effects of these difference on learning is one of the contributions of this work.

With all these caveats in mind, it is perhaps surprising that at a high level, our findings are consistent with the spirit of Kreps et al. [1982]: i.e. that players are not “truly” altruistic, but rather cooperate for strategic reasons. Our conclusion is based on two observations. First, in the repeated public goods game, which of the three experiments most closely resembles the setup for [Kreps et al. 1982], we do indeed see declining rates of cooperation over the long run. Second, in all three experiments, we see increasingly sharp “end-game” effects, meaning that experienced players switch from high to low levels of cooperation more rapidly than inexperienced players, spending less time at intermediate levels. This finding is also consistent with the hypothesis that cooperative players are cooperating at least in part to exploit their position rather than simply for the sake of cooperating.

Attached to these very general and consistent observations, however, are three further caveats. First, although we find diminishing rates of cooperation in the repeated public goods game, the rate of learning is rather slow; specifically, we find that initial rates of cooperation decline by only 50% over 110 games of 10 rounds each. Extrapolating from this finding, moreover, we conjecture that convergence to Nash would require in excess of 200 games of 10 rounds each. Second, although we see sharper end-game effects in all three instances, we do not consistently find lower overall rates of cooperation. Quite to the contrary, in fact, in both the dynamic networks and collaborative search experiments, overall rates of cooperation increase with experience, mostly on account of increases in cooperative behavior in the early rounds (i.e. prior to the end-game effect). Finally, we reiterate that although the data derived from randomized controlled lab experiments these experiments were not designed with long-run learning in mind, hence we must treat the data as observational as opposed to experimental.

³Reflecting a similar assumption, experimental studies of cooperation typically go to considerable lengths to ensure that all subjects have the same experience: they are drawn from a relatively homogeneous population (usually college students), receive the same training, and play the same number of games.

Our findings therefore, should be regarded less as conclusions than as hypotheses to be tested by future, appropriately designed experiments.

The remainder of this paper proceeds as follows. In the next section, we review related work and highlight the major differences between this work and our own. Then, in Section 3, we describe the Investment Game, a repeated public goods game, and show (a) that long-run play converges to unilateral defection, (b) that the so-called end-game effect creeps forward, but that (c) both effects happen slowly. Next, in Sections 4 and 5 we examine the same effects in two other games of cooperation: the “Social Networking Game,” a variant of the repeated Prisoner’s Dilemma game in which players can make and break ties with each other, and “Wildcat Wells,” a game of collaborative search. Although we observe long-run learning effects in both games, we also identify some notable differences arising from the presence of rewiring and re-framing respectively. In particular, we see initial cooperation increasing over time, but end-game effects becoming increasingly sharp. Finally, in section 6 we conclude with some remarks about the theoretical and methodological implications of our work, as well some suggestions for future experiments.

2. RELATED WORK

As noted above, the theoretical literature on learning makes two predictions: one, that with experience, players should “learn” to be less cooperative; and the other that they should learn to become more so. Experimental evidence has also generated ambiguous results with respect the learning hypothesis. Selten and Stoecker [1986] studied a series of 25 prisoner’s dilemma games of 10 rounds each, and found that defection tended to creep earlier over time, in support of the rationality hypothesis. Andreoni and Miller [1993], however, found precisely the opposite result: over the course of 20 PD games, also of 10 rounds each, they found that defection tended to occur later with experience, consistent with altruism. Meanwhile, earlier work by Andreoni [1988] found no evidence to support either direction, and other experiments that use experienced players generate indistinguishable results from those that use naive players [Isaac and Walker 1988]. Precisely how or if learning impacts play in repeated games therefore remains an unresolved question.

In light of this related experimental work and the related theoretical work described in Section 1, our paper makes four main contributions to the literature on long-run learning. First, because we have observations of individuals playing upwards of 100 games, we are able to study learning effects over a much longer timescale than in previous work. Second, in contrast previous studies, which focus on only one very specific game of cooperation—the iterated PD—we consider three rather distinct games of cooperation, hence we are in a position to show that learning patterns depend on certain features of the game in ways that not anticipated by any existing theory. Third, because we have much more data than previous experiments, we are able to analyze not only the main effects of experience (on average payoff and contribution) but also the interaction effects between experience and the game-round, allowing us to examine the dynamics of so-called endgame effects in more detail than previously possible. And finally, whereas previous experiments have studied populations in which all players have the same experience, we study heterogeneous populations—a design feature that more closely resembles real-world “games” of cooperation, in which individuals very likely differ in age and experience.

3. COOPERATION ON STATIC NETWORKS

Suri and Watts [2011] (henceforth SW) conducted a variant of a linear public goods game [Ledyard 1995], a game of cooperation that is widely studied in laboratory settings. Each game comprised 10 rounds, where in each round each participant i was

allocated an endowment of $e = 10$ points, and was required to contribute $0 \leq c_i \leq e$ points to a common pool. Players' payoffs were given by $\pi_i = e_i - c_i + \frac{a}{k+1} \sum_{j \in \Gamma(i)} c_j$, where $\Gamma(i)$ was defined to include i and all its network neighbors, and k was the vertex degree (all nodes in all networks had the same degree). Therefore, i 's contributions were, in effect, divided equally among the edges of the graph that are incident on i , where payoffs are correspondingly summed over i 's edges. From this payoff function it is easy to show that when $1 < a < n$, players face a social dilemma in that all players contributing the maximum amount maximizes social welfare, but individually players are best off if they contribute nothing, thereby free-riding on the contributions of others.

3.1. Experiment and Data

Initially, SW performed 70 preliminary experiments comprising groups of 4, 8 or 16 players each, intended to familiarize players with the game and recruit them to the panel. After that, SW chose networks that spanned a wide range of possible structures between a collection of four disconnected cliques at one extreme, and a regular random graph at the other, where all networks comprised $n = 24$ players, each with constant vertex degree $k = 5$. Next, SW conducted 69 experiments using these networks where all players were humans. After that, SW conducted an additional 39 experiments where at least one, and most often four, dummy players designed to play in prescribed manner. Thus, a total of 108 networked experiments with $N = 24$ were conducted over a period of 1–2 months.⁴ Finally, roughly six months after completing their networked experiments, SW also studied two variants on the networked game, in which the rules of the game were identical but where the players' view of their networks differed either by exposing (a) links between neighbors, or (b) links to the full two-hop neighborhood. These comprised 28 more experiments. Overall, SW conducted a total of 206 experiments over the course of roughly one year. A total of 315 unique individuals participated in these experiments, where as Fig. 1 (bottom right) shows the majority of subjects played only once, while a small number participated in more than 100. Fig. 1 (bottom left) shows per-contributions averaged over all all-human experiments. As can be seen, average contributions start around 6 and then decline roughly linearly, consistent with the canonical finding mentioned previously [Ledyard 1995].

3.2. Overall cooperation

To motivate our analysis of learning effects, Fig. 2 shows the raw data for (a) payoffs and (b) per-round contributions respectively for the Investment Game as a function of experience⁵. Fig. 2 suggests two main results: first, that payoffs in general increase with experience; and second, that contributions generally decline up until some point, at which they increase again. Although the raw data provide a useful starting point, there are several reasons to mistrust these visual impressions. First, as noted above SW conducted several sets of experiments in succession, including some (especially the dummy experiments, but also to a smaller extent the changes of view) with significant treatment effects. Critically, while SW were careful to randomize assignment of players to treatments within each set of experiments, no such randomization was applied across sets. Hence some of the apparent variations with experience are in fact

⁴SW only included experiments where over 90% of the overall contributions and at least 50% of each individual contribution were made by human players as opposed to a default action due to a player dropping out. Since we are interested in learning effects of the human player we included all experiments in this work.

⁵By experience, we mean the experience that player p had at the time of playing experiment q . Thus experience is to be distinguished from *maximum experience* which is the experience that player p accumulates over their entire participation in the Investment Game.

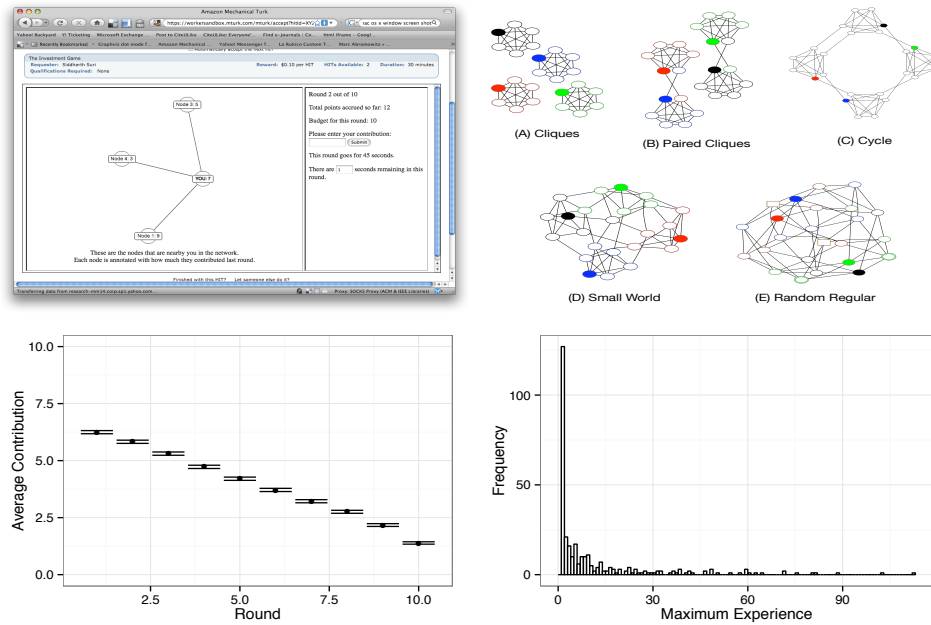


Fig. 1. Clockwise from top left: screen shot of Investment Game; network topologies (with dummy players indicated in solid color); histogram of player experience; and average contribution by round




Fig. 2. Raw data for Investment Game: average cumulative payoff by experience (left); and average per-round contribution by experience (right)

treatment effects. Second, the extreme heterogeneity of experience apparent in Fig. 1 (bottom right) results in many more observations of some players (those with experience) than of others, hence estimates of individual learning must disambiguate within-individual learning from selection effects—namely that players who voluntarily play many games may be systematically different from those who play few games, either because they are different at the outset, or because success in the game itself affects attrition. Finally, individual experiments varied considerably in terms of initial levels of cooperation. Because these random initial differences are known to persist throughout the game, and because they may vary in ways that correlate with experience, we must separately account for experiment-level effects.

Table I. Coefficients of Models fit to Investment Game Data

	Payoff	Contribution	Contribution (Max Experience)	Contribution (by round)
(Intercept)	75.84*** [45.25; 106.46]	2.59* [0.49; 4.72]	2.90** [0.84; 4.99]	-1.45 [-3.91; 1.01]
experience	0.15** [0.05; 0.25]	-0.01*** [-0.02; -0.01]		-0.02*** [-0.02; -0.01]
max experience			0.00 [-0.02; 0.01]	
round				-0.37*** [-0.35; -0.39]
experience:round				-0.0018*** [-0.0012; -0.0023]

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

To address these issues, we adopt a multi-level modeling approach⁶ [Gelman and Hill 2007]. Specifically, we fit models of the following form:

$$y_i = \alpha_{p[i]} + \gamma_{q[i]} + \beta_1 x_i + \beta_2 g_i + \beta_3 c_i + \beta_4 v_i + \epsilon_i \quad (1)$$

where i is an observation of a particular player who played in a particular game. The dependent variable y_i represents either payoff (cumulative over the entire game) or average contribution. Following this notation, $p[i]$ extracts the player from i and $q[i]$ extracts the game from i . Thus $\alpha_{p[i]}$ is a random effect for the players, $\gamma_{q[i]}$ is a random effect for the games, x_i is the experience (or the maximum experience) of the player in that game, g_i is the graph structure, c_i is the treatment condition (e.g., number of dummies), and v_i is the graph distance from the user that is visible to the user. The coefficients for key variables can be seen in Table I; the best-fitting parameters for the full models can be found in the Appendix.

Coefficients for multilevel models can be hard to interpret directly, both on account of the complexity of the model and also differences in scales corresponding to the covariates. Having estimated Eqn. 1, therefore, we can use the fitted model to generate model-adjusted values for payoffs and contributions. That is, for each observation in our data, we set the fixed effects to their observed value, and set the random effects to zero. This procedure preserves the relative frequency of the different treatments, but smooths out some of the variance arising from individual player idiosyncrasies and game effects. Fig. 3 shows these model-adjusted fits, which are analogous to those in Fig. 2, but are noticeably smoother⁷. Nevertheless, the overall impression from the raw data remains: payoffs increase with experience and contributions decrease; thus players are clearly “learning” both in the sense that they are performing better than inexperienced players, and also that their contributions are decreasing, consistent with “rational” interpretation of the incomplete information hypothesis.

Although our inclusion of a random effect for individual players should account for biases introduced by selection effects (e.g. that inherently less-generous players are more successful and also are more motivated to play many games) we also address the selection issue directly by fitting the analogous model to Eqn. 1, but with maximum experience (i.e. the experience that player i in experiment j eventually attains) replacing experience as the fixed effect. Interestingly, column 3 of Table I shows that the coefficient for maximum experience is close to zero and not significant, indicating that selection is not a significant effect.

⁶Multilevel models are also known as hierarchical linear models, mixed models, or random-effects models

⁷We note that our simulation procedure preserves the order in which the various values of the fixed effects appear in the raw data, hence some of the fluctuations associated with distinct sets of experiments remain in the model-adjusted fits.




Fig. 3. Model-adjusted data for average cumulative payoff as a function of experience (left), and average per-round contribution as function of experience (right)

3.3. Endgame effect

Aside from the effects of learning on average levels of cooperation, we note that Kreps et al. [1982] also speculated that learning effects may manifest themselves differently at different stages of the game. Specifically, Kreps et al. [1982] noted an “endgame effect” in which cooperation ceases as the end of the game approaches—behavior that is suggestive of strategic cooperation rather than cooperation arising out of “true” altruism. As Selten and Stoecker [1986] also noted, if players are in fact rational, one would expect this endgame effect to creep forward over time, as players iteratively attempt to preempt one other in defecting, leading eventually to total unraveling of cooperation just as predicted by the naive theory.

To study end-game effects, we fit the following multilevel model:

$$y_i = \alpha_{p[i]} + \gamma_{q[i]} + \beta_1 x_i + \beta_2 r_i + \beta_3 x_i r_i + \beta_4 g_i + \beta_5 c_i + \beta_6 v_i + \epsilon_i \quad (2)$$

where y_i is an indicator variable for whether the contribution of player $p[i]$ in round r of game $q[i]$ was greater than 5, and r_i is a fixed effect for round. Note that unlike in Eqn. 1, here we fit a logistic model, where we first dichotomize contribution as greater than or less than 5 points⁸.

Column 4 of Table I shows that the main effects of experience and round are both negative and highly significant. More importantly, however, the interaction effect between experience and round is negative and also highly significant, indicating that the decrease in cooperation with experience is more pronounced in later rounds than in earlier, consistent with the endgame effect creeping forward. Fig. 4 shows this interaction effect in two ways: first, Fig. 4 (left) shows model-adjusted per-round contribution as a function of experience, broken down by round; and second, Fig. 4 (right) shows model-adjusted per-round contribution as a function of round, broken down by experience. In particular, Fig. 4 (right) shows both that initial cooperation decreases with experience, and also that cooperation decreases more rapidly as the game progresses, corresponding to an encroaching endgame effect.

To sum up, our analysis supports the view advanced by Kreps et al. [1982] that players in repeated games of cooperation are in fact rational, and are cooperating at least in part because they believe that other players may not be. In this view, as the players gain experience they become increasingly aware of the other players’ “true” nature and

⁸We use a logistic model in large part because both subsequent experiments involve binary outcome variables, hence for the per-round analysis logistic models are required. However, a logistic model is also not inappropriate for the Investment Game as the vast majority of contributions are concentrated around either 10 or 0, with very little mass in the middle [Suri and Watts 2011], hence converting the real-valued contribution into a binary outcome results in very little loss of detail.




Fig. 4. (Best viewed in color) Model-adjusted data for per-round contribution as a function of experience broken down by round (left), and as a function of round broken down by experience (right)

respond (a) by cooperating less overall, and (b) by reducing their contributions more rapidly as the game progresses, as we indeed see⁹. As striking as these effects appear, however, we also note that the timescales involved are rather large: for example, the decrease in initial cooperation visible in Fig. 4 (left) from approximately 0.6 to approximately 0.3 takes place over 110 games. Extrapolating our model fits, we estimate a lower bound on the number of games required for cooperation to disappear entirely at over 200 games¹⁰. Cooperation, in other words, may indeed disappear in the long run, but the long run is very long. Moreover, as we explain in the next two sections, even slight changes to the structure of the game can prolong cooperation indefinitely.

4. COOPERATION IN DYNAMIC NETWORKS

Wang, Suri, and Watts [2012] conducted a series of online human subjects experiments in which groups of 24 participants played an iterated Prisoner’s Dilemma (PD), where in addition to choosing their action each round—cooperate or defect—they also were given the opportunity to update their interaction partners at some specified rate, which was varied across experimental conditions. All games comprised 12 “strategy update” rounds during which every player could update their strategy: cooperate (C) or defect (D). Consistent with standard PD conditions, a cooperator received R points when interacting with another cooperator, and S points when interacting with a defector, while a defector received T points when interacting with a cooperator and P points when interacting with another defector, where $T > R > P > S$ and $T + S < 2R$. In addition, after every r strategy update rounds, players entered a “partner updating” turn in which they were permitted to make up to k partner updates. By adjusting r and k Wang et al. [2012] explored a wide range of relative updating rates, from one opportunity every several strategy update rounds to several opportunities every round. A single update comprised either severing a link with an existing partner or proposing

⁹Our empirical observations are broadly consistent with at least two other explanations. First, players could all be conditional cooperators who do not so much learn about each other as simply update their actions via some variant of reinforcement learning, where end-game effects encroach over time as players seek to avoid being exploited (similar to the the hypothesis of Selten and Stoecker [1986]). And second, it could be that all players start as conditional cooperators, but that a small fraction learn to play rationally; the effect is therefore driven entirely by this small fraction who in turn drive down the contributions of the majority. Although our analysis does not differentiate between the rational cooperation hypothesis and either of these two alternatives, it may be possible to do so with individual-level modeling or with better-designed experiments. We are grateful to two anonymous referees for raising these alternatives.

¹⁰We acknowledge that extrapolation out of sample is an inherently unreliable procedure. In particular, the relationship between initial contribution and experience is likely to become increasingly nonlinear as it approaches zero. Hence our estimate is at best a loose lower bound with unknown uncertainty. Nevertheless, it suffices to make the point that cooperation likely persists for hundreds of games.

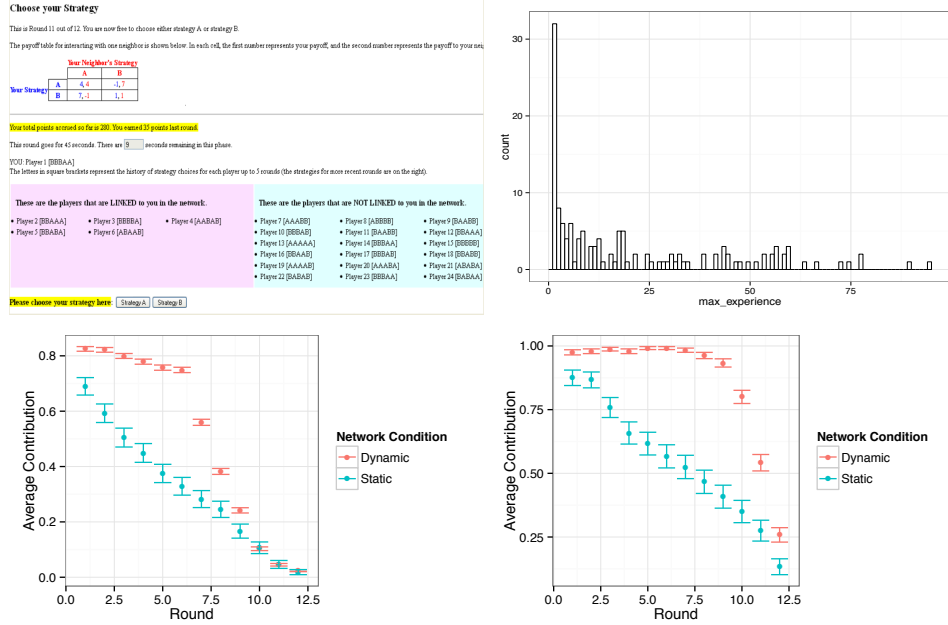


Fig. 5. (Best viewed in color) Clockwise from top left: screen shot of Social Networking Game; histogram of player experience; probability of cooperating by round for payoffs 2; probability of cooperating by round for payoffs 1.

a link to a new partner, where importantly, players could choose the partner to either sever or propose a link to. After each partner updating turn was completed, the network of partners was updated to reflect severed and accepted links, and a new strategy update round commenced.

4.1. Experiment and Data

Wang et al. [2012] conducted two sets of experiments in succession totaling 109 experiments with 114 unique players. The first set of experiments used payoffs $T_1 = 7, R_1 = 4, P_1 = 1, S_1 = -1$ (hereafter payoffs 1); and the second used payoffs $T_2 = 7, R_2 = 4, P_2 = -1, S_2 = -5$ (hereafter payoffs 2). Although both sets of payoffs satisfy the conditions for a PD ($T > R > P > S$ and $T + S < 2R$), they differ with respect to the relative cost and benefit of making and breaking ties. Specifically, a cooperator facing the choice of breaking a tie with a currently defecting partner and forming a new tie with another cooperator will prefer the latter in payoffs 1 and the former in payoffs 2. As a consequence, Wang et al. [2012] found that in payoffs 1, players tended to retain ties with defectors even though these ties were costly, preferring to use their updates to create additional ties with other cooperators, whereas in payoffs 2 they were more likely to punish defecting partners by severing ties. As Fig. 5 (bottom row) shows, rewiring led to large and significant increases in cooperation (relative to the static controls) for both sets of payoffs. However, payoffs 2 not only led to significantly higher cooperation levels initially, but also maintained cooperation for longer.

4.2. Overall cooperation

Fig. 6 shows the raw data for average player payoff and probability of cooperation (i.e. corresponding to Fig. 2 for the Investment Game), where we have separated out




Fig. 6. Raw data for Social Networking Game with payoffs 1 (top row) and payoffs 2 (bottom row). Left column represents average cumulative payoff and right column is average probability of cooperating, both as function of experience.

payoffs 1 and 2—treating them, in effect, as two distinct experiments¹¹. For similar reasons, we also omit the data from the control conditions, fitting only the treatments in which rewiring occurs. Consistent with Wang et al. [2012], Fig. 6 shows that both cooperation and payoffs are higher in the dynamic than in the static conditions, and that the effects are larger for payoffs 2 than payoffs 1.

Unlike for the Investment Game, the raw data does not show any clear trends as a function of experience. As before, however, the raw data is potentially confounded by treatment effects, individual differences, and variability from game to game. To account for these confounding effects, we again fit multilevel models of the form:

$$y_i = \alpha_{p[i]} + \gamma_{q[i]} + \beta_1 x_i + \beta_2 g_i + \beta_3 s_i + \beta_4 w_i + \beta_5 s_i w_i + \epsilon_i \quad (3)$$

for payoffs 1, where i is an observation of a particular player who played in a particular game, and $p[i]$ extracts the player from i and $q[i]$ extracts the game from i . The dependent variable y_i represents either payoff (cumulative over the entire game) or probability of cooperating for player $p[i]$ in game $q[i]$, $\alpha_{p[i]}$ is a random effect for the players, $\gamma_{q[i]}$ is a random effect for the individual game, x_i is the experience (or

¹¹Fig. 5 shows the behavior of players differed dramatically between payoffs 1 and 2, which were conducted in non-overlapping batches; hence it makes sense to treat them as distinct experiments. Because there was overlap between the panels for both games, however, treating them separately introduces an additional complication for the second payoffs—namely that “experience” is no longer uniquely defined. One possibility is to measure total experience (i.e. in both payoffs 1 and 2), but this choice effectively treats n games of payoffs 1 and m games of payoffs 2 as identical to $n + m$ games of payoff 2, which, given the differences between the two payoff schemes, may be misleading. A second possibility is to ignore players’ experience with payoffs 1, effectively counting them as “fresh” when they play their first game with payoffs 2. This assumption, however, would conflate truly inexperienced players with players who have extensive experience with a closely related game—also likely misleading. We adopt the following compromise: We set experience with payoff 2 as the fixed effect of interest, but also include a random effect for experience in payoff 1.

Table II. Coefficients of Models Fit to Social Networking Game: Payoffs 1

	Payoff	Coop. Likelihood	Coop. Likelihood (Max Experience)	Coop. Likelihood (by round)
(Intercept)	274.28*** [247.89; 300.90]	0.61*** [0.56; 0.66]	0.55*** [0.48; 0.62]	4.25*** [3.72; 4.77]
experience	0.42* [0.00; 0.80]	0.00*** [0.00; 0.00]		0.04*** [0.03; 0.05]
max experience			0.00 [0.00; 0.00]	
round				-0.50*** [-0.53; -0.48]
experience:round				-0.0011*** [-0.001; -0.0012]

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

the maximum experience) of player $p[i]$ in game $q[i]$, g_i represents the different graph structures, s_i represents the frequency of frequency of partner update turns, and w_i represents the number of updates allowed per partner update turn.

For payoffs 2, all initial graph structures were the same, and all players had a rewiring opportunity every round, so we fit models of the form:

$$y_i = \alpha_{p[i]} + \gamma_{q[i]} + \delta_{m[i]} + \beta_1 x_i + \beta_2 w_i + \epsilon_i \quad (4)$$

For payoffs 2, some of the players had previously played with payoffs 1. So here i is an observation that a particular player, $p[i]$, played in a particular game, $q[i]$, with payoffs 2 after playing a particular number of games, $m[i]$, with payoffs 1. We included a random effect $\delta_{m[i]}$ for the number of games the individual had played with payoffs 1. Table II shows the estimated coefficients, and Fig. 7 shows the corresponding model-adjusted data for payoff and cooperation, for both sets of payoffs. Payoffs 1 (top row) are similar to the analogous results for the Investment Game (Fig. 3), in that payoffs show a gradual increase with learning while cooperation slowly decreases. Given the differences between the two games, and the overall higher rates of cooperation in the Social Networking games, this similarity is somewhat surprising, and suggests that learning effects in games of cooperation generalize somewhat beyond the specifics of any one game. The results for Payoffs 2 (bottom row), however, caution that any such generalizability is limited. Not only do payoffs 2 lead to higher overall rates of cooperation than payoffs 1, that is, but the direction of the learning effect is reversed: players are clearly becoming more cooperative with experience, not less.

4.3. End of Game Effects

Another striking difference between the Social Networking Game and the Investment Game, clearly apparent in Fig. 5 (bottom row), is that the endgame effect is much sharper in the Social Networking Game. Moreover, payoffs 1 and 2 also differ substantially in the behavior during round 6 and beyond. Payoffs 1 exhibit a defection cascade starting in round 6 whereas payoffs 2 exhibit high levels of cooperation at least round 10 when the end game effect occurs. To investigate the evolution of these effects in the long run, we again fit multilevel logistic regression models of the following form:

$$y_i = \alpha_{p[i]} + \gamma_{q[i]} + \beta_1 x_i + \beta_2 r_i + \beta_3 x_i r_i + \beta_4 g_i + \beta_5 s_i + \beta_6 w_i + \beta_7 s_i w_i + \epsilon_i \quad (5)$$

for payoffs 1 where y_i is an indicator variable for whether player $p[i]$ cooperated in round r of game $q[i]$, and r_i is a fixed effect for round. Again, this model had to be modified for payoffs 2, taking the form:

$$y_i = \alpha_{p[i]} + \gamma_{q[i]} + \delta_{m[i]} + \beta_1 x_i + \beta_2 r_i + \beta_3 x_i r_i + \beta_4 w_i + \epsilon_i \quad (6)$$

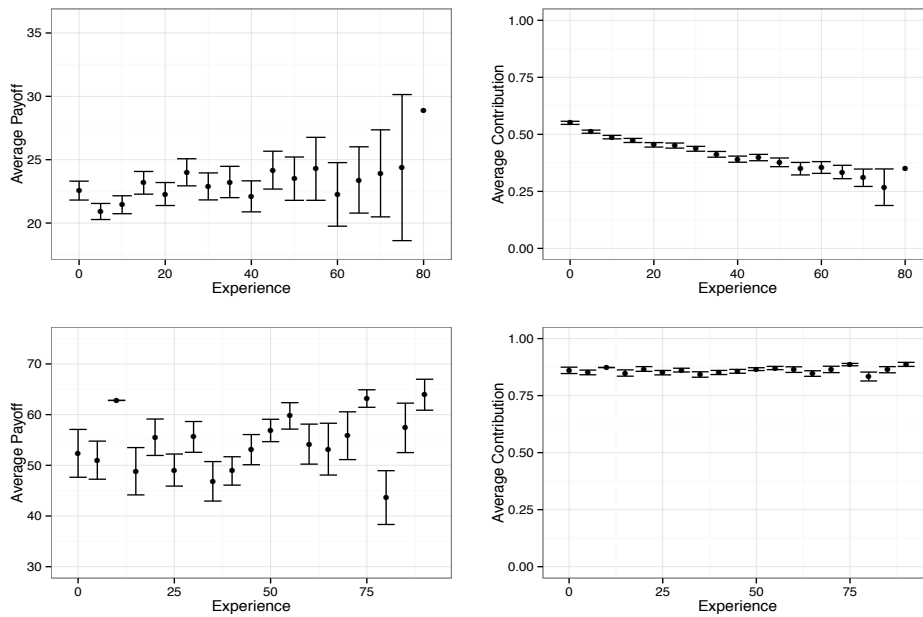


Fig. 7. Model-adjusted data for social networking game with payoffs 1 (top row) and payoffs 2 (bottom row). Left column represents average cumulative payoff and right column is average probability of cooperating, both as function of experience.

Column 4 of Tables II and III shows the estimated coefficients for Eqn. 5 and 6 respectively. Fig. 8 shows the corresponding model-adjusted data, and reveals both similarities and differences with Fig. 4. On the one hand, both display strong interaction effects between experience and round, consistent with the endgame effect increasing in strength as players gain experience. On the other hand, whereas in the Investment Game cooperation decreased with experience at all stages of the game, here we find initial cooperation increases with experience in both payoffs 1 and 2. Players in other words, appear to be learning to cooperate more at the start of the game and then exploiting the ties they gain as cooperators by switching to defection as the game progresses, where the switch becomes increasingly sharp with experience. Finally, the direction in which the endgame effect moves differs depending on the payoffs: for payoffs 1, it clearly creeps forward, consistent with the Investment Game, converging on round 5; but for payoffs 2 it recedes with experience, converging on round 10.

5. COLLABORATIVE LEARNING EXPERIMENT

Mason and Watts [2012] (hereafter MW) conducted an experiment, “Wildcat Wells,” in which players were tasked with exploring a realistic-looking two-dimensional desert map in search of hidden “oil fields.” The players had 15 rounds to explore the map, either by entering grid coordinates by hand, or by clicking directly on the map (see Fig. 9 (top left)). On each round after the first, players were shown the history of their searched locations and payoffs, as well as the history of searched locations and payoffs of three collaborators. In choosing their next location, therefore, players repeatedly faced a choice between copying the best current score of their neighbors or exploring new terrain on their own. MW showed that the decision to explore was beneficial to the collective performance but copying improved the score of the copier, hence players effectively faced a repeated social dilemma in the explore-copy decision [Mason and

Table III. Coefficients of Models Fit to Social Networking Game: Payoffs 2

	Payoff	Coop. Likelihood	Coop. Likelihood (Max Experience)	Coop. Likelihood (by round)
(Intercept)	372.44*** [291.21; 453.24]	0.80*** [0.73; 0.87]	0.79*** [0.72; 0.87]	3.3*** [1.68; 4.92]
experience	0.61 [-0.81; 2.06]	0.00 [0.00; 0.00]	0.00 [0.00; 0.00]	0.19*** [0.14; 0.25]
max experience				
round				-0.3*** [-0.16; -0.45]
experience:round				-0.019*** [-0.014; -0.024]

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

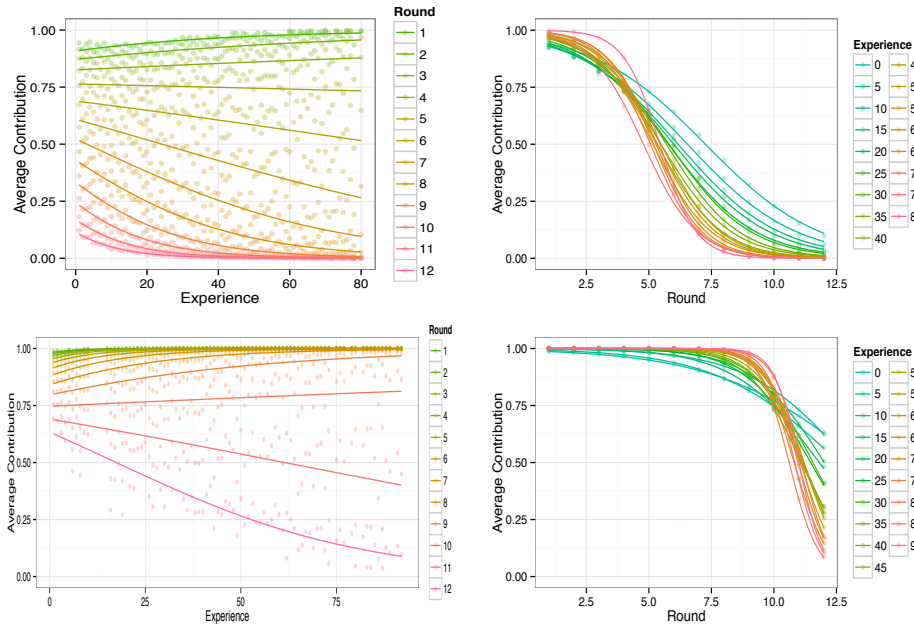


Fig. 8. (Best viewed in color) Social Networking Game payoffs 1 (top row) and payoffs 2 (bottom row): model adjusted fits for per-round contribution as a function of experience, broken down by round (left column), and as a function of round, broken down by experience (right column).

Watts 2012]¹². Supporting this interpretation, Fig. 9 (bottom left) shows that, players mostly explored new locations in early rounds while in later rounds they were more likely to copy their most successful neighbor.

5.1. Experiment and Data

Each experimental session comprised 8 games corresponding to each of the network topologies, so players experienced each topology exactly once in random order. Players

¹²The fitness landscape comprised two components: a main peak and noisy background. Once a neighbor had found the main peak, an experienced player might recognize that further exploration was pointless, hence removing the dilemma. For this reason, we restrict attention here to rounds on which the focal player's neighbors had not found the peak.

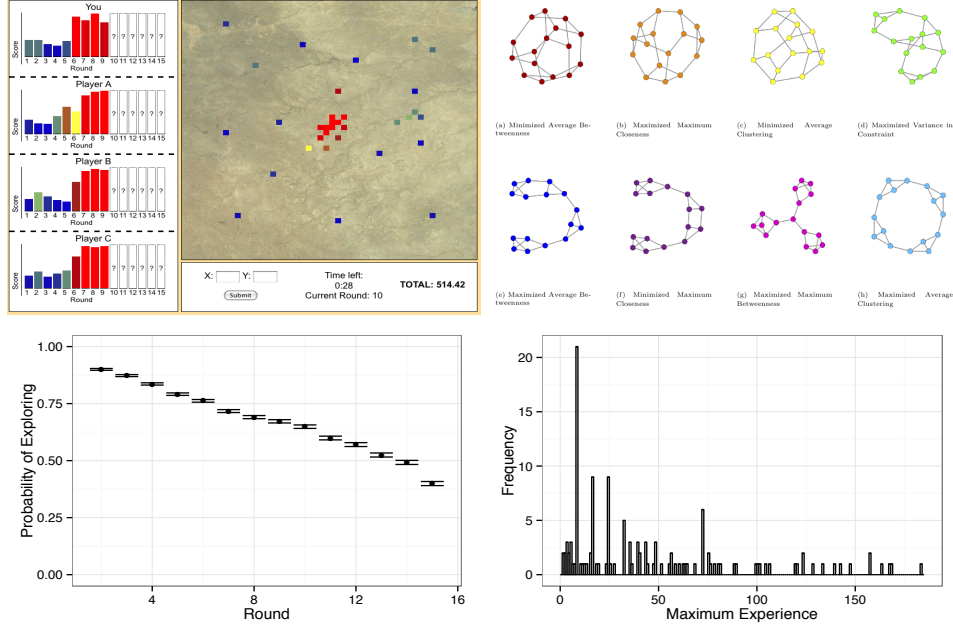


Fig. 9. Clockwise from top left: screen shot of Wildcat Wells; networks studied in the experiment; histogram of player experience; probability of exploring (cooperating) as a function of round.

were randomly assigned to unique positions in one of 8 network topologies, where each player’s collaborators for that game were his or her immediate neighbors in that network. All networks comprised $n = 16$ nodes, each with $k = 3$ neighbors, but differed with respect to four commonly-studied metrics that have been posited to affect information flow in networks [Mason and Watts 2012]. In this manner, MW conducted 232 networked games over 29 sessions, as well as a series of 24 baseline experiments, in which groups of 16 individuals searched the same landscape independently (i.e. with no network neighbors and no sharing of information), resulting in a total of 256 experiments with 130 unique individuals.

5.2. Average payoffs and cooperation

As with the previous two experiments, we first show the raw data (Fig. 10) and then show the model-adjusted data (Fig. 11) based on coefficients estimated from the following multilevel model:

$$y_i = \alpha_{p[i]} + \gamma_{q[i]} + \delta_{m[i]} + \beta_1 x_i + \beta_2 g_i + \epsilon_i \quad (7)$$

where i is an observation of a particular player who played a particular trial of a particular game, and $p[i]$ extracts the player from i , $m[i]$ extracts the trial from i , and $q[i]$ extracts the game from i . As before the dependent variable y_i represents either estimated payoff (cumulative over the entire game) or the estimated probability of exploring for player $p[i]$ in trial $m[i]$ of game $q[i]$; $\alpha_{p[i]}$ is a random effect for the players, $\delta_{m[i]}$ is a random effect for the trial within the experiment, $\gamma_{q[i]}$ is a random effect for the individual game, x is the experience of player, and g represents the graph structure. The coefficients estimated from Eqn. 7 are given in the first two columns of Table IV.

Focusing on Fig. 11 (the model-adjusted fits), we see a pattern that is most similar to payoffs 2 of the Social Networking Game, in that payoffs decrease with experience and cooperation (exploration) increases. The magnitude of the coefficients, however, is




Fig. 10. Wildcat Wells, raw data for average cumulative payoff as a function of experience (left), average per-round probability of exploring as function of experience (right)




Fig. 11. Wildcat Wells, model adjusted fits for average cumulative payoff as a function of experience (left), average per-round probability of exploring as function of experience (right)

Table IV. Coefficients of Model Fit to Wildcat Wells Data

	Payoff	Prob. Exploring	Prob. Exploring (Max Experience)	Prob. Exploring (by round)
(Intercept)	512.33*** [482.84; 542.00]	0.72*** [0.68; 0.76]	0.72*** [0.67; 0.78]	2.63*** [2.31; 2.95]
experience	-0.08 [-0.31; 0.15]	0.00 [0.00; 0.00]		0.014*** [0.011; 0.016]
max experience			0.00 [0.00; 0.00]	
round				-0.17*** [-0.16; -0.18]
experience:round				-0.0011*** [-0.001; -0.0013]

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

consistently smaller than in either of the previous experiments, suggesting that the extra complexity of Wildcat Wells made it a more difficult game to “learn.”

5.3. Endgame Effect

Considering now the endgame effects, we fit multilevel logistic models of the form

$$y_i = \alpha_{p[i]} + \gamma_{q[i]} + \delta_{m[i]} + \beta_1 x_i + \beta_2 r_i + \beta_3 x_i r_i + \beta_4 g_i + \epsilon_i \quad (8)$$

where y_i is an indicator variable for whether the player $p[i]$ in round r of trial $m[i]$ in game $q[i]$ copied one of her neighbors, and r is a fixed effect for round.

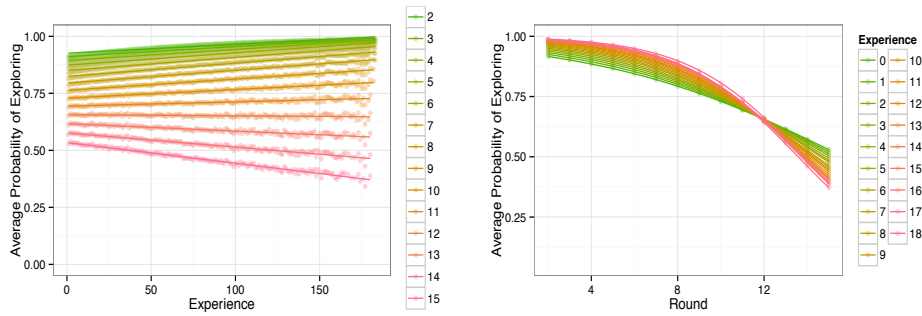


Fig. 12. (Best viewed in color) Wildcat Wells model adjusted fits for per-round probability of exploring broken as a function of experience, broken down by round (left); and as a function of round, broken down by experience (right)

The coefficients for Eqn. 8 are given in column 4 of Table IV and the model adjusted data for per-round probability of exploring as a function of experience (broken down by round) and round (broken down by experience) are shown in Fig. 12. Similar to Fig. 11, Fig. 12 suggests that Wildcat Wells is more similar to the Social Networking Game than to the Investment Game, in that initial cooperation is increasing while the endgame effect becomes increasingly sharp, where the shape of the curves are roughly intermediate between payoffs 1 and 2 of the Social Networking Game.

These results are somewhat surprising, as the tradeoff between exploring and copying in Wildcat Wells seems at first glance more similar to the Investment Game than to the Social Networking Game, where the capability to rewire links can be thought of as a form of punishment for defection. No such punishment mechanism is available to players in Wildcat Wells, which is therefore more like a classic repeated public goods game, aka the Investment Game. Clearly, however, Wildcat Wells is not framed explicitly as a game of cooperation, nor is the social dilemma inherent in payoffs necessarily obvious to the players. Even though they are, in effect, playing a game cooperation, that is, they may not realize it, and this obfuscation may affect their behavior. Although not surprising, we note that real-world “games” of cooperation are also unlikely to be presented in terms of explicit payoffs, as they are in standard experimental designs, hence the learning effects—or lack thereof—in the Wildcat Wells game may more closely resemble real life.

6. DISCUSSION

To summarize, we have analyzed data from three distinct cooperation games with players with heterogeneous experience levels: a linear public goods game, a multi-player prisoner’s dilemma with partner updating and a collaborative explore/exploit game. Despite the variety in game type we see learning effects across all three games. Specifically, we find that as people gain experience, they behave in a more extreme manner, in the sense that where the contribution curves are concave the curve becomes more concave, and where they are convex they become more convex. Furthermore, across all three games there is an interaction effect between experience and round. As people get more experience, they learn to behave more selfishly in the last round; thus even where we see initial increases in cooperation, as we do in the Social Networking and Wildcat games, it is likely strategic. Overall our findings are therefore consistent with the “rational cooperation” hypothesis of Kreps et al. [1982], according to which players eventually learn to act selfishly. The Investment Game analysis shows this result most clearly, with players clearly converging towards the equilibrium strategy predicted by the standard backward induction argument. Even for the Investment Game, however,

we found that it would take a long time, at least 200 games, for players to actually reach the equilibrium. An interesting challenge for future theoretical work, would be to explain not only the equilibrium strategy, but also the timescale required for players to converge to it. Our results for the Social Networking Game and Wildcat Wells also present interesting challenges for theory: both displayed significant learning effects, but to our knowledge there is not a well developed theory of learning in these settings. Explaining the direction and type of learning effects, and how they depend on particular details of the games, stands as a challenge for future theoretical work.

REFERENCES

- ANDREONI, J. 1988. Why free ride?: Strategies and learning in public goods experiments. *Journal of public Economics* 37, 3, 291–304.
- ANDREONI, J. 1995. Cooperation in public-goods experiments: kindness or confusion? *The American Economic Review*, 891–904.
- ANDREONI, J. AND MILLER, J. H. 1993. Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence. *The Economic Journal* 103, 418, 570–585.
- GELMAN, A. AND HILL, J. 2007. *Data analysis using regression and multi-level/hierarchical models*. Vol. Analytical methods for social research. Cambridge University Press, New York.
- HORTON, J. J., RAND, D. G., AND ZECKHAUSER, R. J. 2011. The online laboratory: Conducting experiments in a real labor market. *Experimental Economics* 14, 3, 399–425.
- ISAAC, R. M. AND WALKER, J. M. 1988. Group size effects in public goods provision: The voluntary contributions mechanism. *The Quarterly Journal of Economics* 103, 1, 179–199.
- KREPS, D. M., MILGROM, P., ROBERTS, J., AND WILSON, R. 1982. Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic theory* 27, 2, 245–252.
- LEDYARD, J. 1995. Public goods: A survey of experimental research. In *Handbook of Experimental Economics*, J. H. Hagel and A. E. Roth, Eds. Princeton University Press, Princeton, NJ, 111–194.
- MASON, W. AND SURI, S. 2012. Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods* 44, 1, 1–23.
- MASON, W. AND WATTS, D. J. 2012. Collaborative learning in networks. *Proceedings of the National Academy of Sciences* 109, 3, 764–769.
- OSBORNE, M. J. AND RUBINSTEIN, A. 1994. *A Course in Game Theory*. MIT Press.
- PAOLACCI, G., CHANDLER, J., AND IPEIROTIS, P. 2010. Running experiments on amazon mechanical turk. *Judgment and Decision Making* 5, 5, 411–419.
- SELTEN, R. AND STOECKER, R. 1986. End behavior in sequences of finite prisoner's dilemma supergames a learning theory approach. *Journal of Economic Behavior & Organization* 7, 1, 47–70.
- SURI, S. AND WATTS, D. J. 2011. Cooperation and contagion in web-based, networked public goods experiments. *PLoS One* 6, 3, e16836.
- WANG, J., SURI, S., AND WATTS, D. J. 2012. Cooperation and assortativity with dynamic partner updating. *Proceedings of the National Academy of Sciences* 109, 36, 14363–14368.

Online Appendix to: Long-Run Learning in Games of Cooperation

WINTER MASON, Facebook
SIDDHARTH SURI, Microsoft Research
DUNCAN J. WATTS, Microsoft Research

A. APPENDIX

Tables V-VIII give coefficients for the full models (equations 1-8) given in the text. The extra coefficients refer to the various fixed effects arising from different network topologies and experimental treatments (see descriptions of experiments, main text). For example, in Table V.

- “experience:round” is an interaction effect between experience and round
- rows of the form “ $m k_n$ ” refers to graphs with m cliques of n individuals each
- The “paired cliques”, “cycle”, “small world” and “random regular” topologies are described in Suri and Watts [2011].
- “cube” is a graph where the nodes are the corners of a three dimensional cube. The vertices are connected by edges in the same way as the cube.
- “ $2 k_4$ joined” is like the paired cliques topology of Suri and Watts [2011] except there are only two cliques of four vertices each, as opposed to four cliques of six vertices each.
- “dummy 10:cover” refers to an experimental treatment in which four dummy players were inserted such that every human player was exposed to exactly one dummy. The dummy players contributed their full endowment (10 points) on each round (see [Suri and Watts 2011]).
- “dummy 10:neighbor” refers to an experimental treatment in which four dummy players were inserted such that some human players were exposed to exactly two dummies. The dummy players contributed their full endowment (10 points) on each round. Other human players were not connected to any dummies (see [Suri and Watts 2011]).
- “dummy 10:single” refers to an experimental treatment in which one dummy player was inserted and it contributed its full endowment (10 points) on each round.
- “no dummies:human” refers to an experimental treatment in which all of the players were humans.
- “ego view” refers to a variant in which players could see which of their nearest neighbors were connected.
- “2-hop view” refers to a variant in which players could all nodes and edges within a graph distance of two from themselves.

In Tables VI and VII:

- “ $r = x$ ” refers to the condition in which players experienced a partner updating turn every x strategy update rounds
- “ $k = y$ ” is the condition in which players could rewire y links every partner update turn (see [Wang et al. 2012]).
- “ $r = x : k = y$ ” refers to the interaction effect of rewiring y links every x rounds.

Copyright © 2013 ACM 978-1-4503-2565-3/14/06...\$15.00
DOI 10.1145/http://dx.doi.org/10.1145/2600057.2602892 <http://doi.acm.org/10.1145/http://dx.doi.org/10.1145/2600057.2602892>

In Table VIII “max max betweenness” to “max var constraint” are fixed effects for seven of the eight network topologies (because these are fixed effects, the eighth topology is the reference). See Mason and Watts [2012] for the definition of these networks.

Table V. Coefficients of Models fit to Investment Game Data

	Payoff	Contribution	Contribution (Max Experience)	Contribution (by round)
(Intercept)	75.84*** [45.25; 106.46]	2.59* [0.49; 4.72]	2.90** [0.84; 4.99]	-1.45 [-3.91; 1.01]
experience	0.15** [0.05; 0.25]	-0.01*** [-0.02; -0.01]		-0.02*** [-0.02; -0.01]
max experience			0.00 [-0.02; 0.01]	
round				-0.37*** [-0.35; -0.39]
experience:round				-0.0018*** [-0.0012; -0.0023]
k_8	7.39 [-26.79; 41.66]	-0.09 [-2.43; 2.24]	-0.03 [-2.32; 2.26]	0.32 [-2.21; 2.84]
10 k_6	87.72*** [47.38; 128.04]	2.59 [-0.19; 5.34]	2.39 [-0.31; 5.09]	2.49 [-0.66; 5.64]
11 k_6	113.90*** [77.91; 149.86]	1.58 [-0.91; 4.03]	1.28 [-1.12; 3.68]	1.97 [-0.8; 4.73]
2 k_4	11.90 [-19.58; 43.42]	0.27 [-1.88; 2.42]	0.27 [-1.84; 2.37]	0.47 [-1.87; 2.81]
2 k_4 joined	9.39 [-22.06; 40.93]	-0.46 [-2.61; 1.69]	-0.49 [-2.59; 1.62]	-0.26 [-2.6; 2.09]
4 k_6	44.92** [14.17; 75.65]	-0.18 [-2.30; 1.92]	-0.40 [-2.46; 1.66]	0.23 [-2.09; 2.56]
cycle	44.59** [14.36; 74.76]	-0.06 [-2.16; 2.00]	-0.46 [-2.48; 1.56]	0.67 [-1.61; 2.94]
paired cliques	47.20** [16.51; 77.86]	0.27 [-1.85; 2.37]	-0.07 [-2.13; 1.98]	0.78 [-1.53; 3.1]
small world	45.96** [15.43; 76.46]	-0.19 [-2.30; 1.89]	-0.43 [-2.48; 1.61]	0.24 [-2.06; 2.55]
random regular	46.60** [16.15; 77.02]	0.10 [-2.01; 2.18]	-0.18 [-2.22; 1.86]	0.61 [-1.69; 2.91]
8 k_6	97.05*** [56.35; 137.72]	2.27 [-0.53; 5.04]	2.05 [-0.67; 4.75]	2.22 [-0.94; 5.38]
cube	8.04 [-23.41; 39.60]	-0.48 [-2.63; 1.67]	-0.48 [-2.58; 1.63]	-0.21 [-2.56; 2.14]
k_4	12.42 [-18.08; 42.89]	1.05 [-1.06; 3.14]	0.88 [-1.18; 2.94]	1.43 [-0.85; 3.71]
dummy 10:cover	59.14*** [49.55; 68.74]	2.38*** [1.74; 3.03]	2.40*** [1.77; 3.03]	1.84*** [1.11; 2.58]
dummy 10:neighbor	29.76*** [20.53; 39.00]	1.32*** [0.66; 1.94]	0.96** [0.36; 1.54]	1.67*** [0.98; 2.36]
dummy 10:single	8.92 [-8.48; 26.45]	0.35 [-0.82; 1.54]	0.48 [-0.67; 1.63]	-0.18 [-1.54; 1.18]
no dummies:human	28.31*** [20.67; 35.97]	1.50*** [0.98; 2.01]	1.46*** [0.96; 1.96]	1.31*** [0.72; 1.9]
ego view	2.47 [-7.36; 12.28]	0.82* [0.14; 1.48]	0.59 [-0.05; 1.23]	1.086** [0.33; 1.84]
2-hop view	17.28*** [7.58; 26.85]	1.14** [0.37; 1.84]	0.81* [0.10; 1.50]	1.74*** [0.98; 2.5]
AIC	34945.72	15598.36	15609.97	
BIC	35093.51	15746.15	15757.76	
Log Likelihood	-17448.86	-7775.18	-7780.98	
Deviance	34897.72	15550.36	15561.97	
Num. obs.	3491	3491	3491	
Num. groups: turk id	315	315	315	
Num. groups: experiment id	206	206	206	
Variance: turk id.(Intercept)	128.02	4.70	4.72	
Variance: experiment id.(Intercept)	172.44	0.86	0.80	
Variance: Residual	1182.84	3.93	3.96	

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table VI. Coefficients of Models Fit to Social Networking Game: Payoffs 1

	Payoff	Coop. Likelihood	Coop. Likelihood (Max Experience)	Coop. Likelihood (by round)
(Intercept)	274.28*** [247.89; 300.90]	0.61*** [0.56; 0.66]	0.55*** [0.48; 0.62]	4.25*** [3.72; 4.77]
experience	0.42* [0.00; 0.80]	0.00*** [0.00; 0.00]		0.04*** [0.03; 0.05]
max experience			0.00 [0.00; 0.00]	
round				-0.50*** [-0.53; -0.48]
experience:round				-0.0011*** [-0.001; -0.0012]
random regular	7.34 [-9.65; 24.37]	0.00 [-0.02; 0.02]	0.00 [-0.03; 0.03]	0.01 [-0.22; 0.24]
$r = 3$	-121.60*** [-155.68; -87.56]	-0.06** [-0.11; -0.02]	-0.05 [-0.11; 0.01]	-0.62** [-1.08; -0.16]
$r = 6$	-198.34*** [-233.38; -163.28]	-0.35*** [-0.40; -0.31]	-0.36*** [-0.42; -0.30]	-3.56*** [-4.04; -3.08]
$k = 3$	194.12*** [157.90; 230.36]	0.00 [-0.05; 0.05]	0.00 [-0.06; 0.06]	-0.02 [-0.51; 0.47]
$k = 5$	264.81*** [230.77; 298.84]	0.08*** [0.04; 0.13]	0.08** [0.03; 0.14]	0.82*** [0.36; 1.29]
$r = 3 : k = 3$	-91.11*** [-142.08; -40.09]	0.07 [0.00; 0.14]	0.06 [-0.03; 0.15]	0.7* [0.01; 1.39]
$r = 6 : k = 3$	-155.42*** [-205.87; -104.98]	0.14*** [0.07; 0.21]	0.14** [0.05; 0.23]	1.5*** [0.81; 2.18]
$r = 3 : k = 5$	-104.55*** [-153.90; -55.15]	-0.03 [-0.10; 0.04]	-0.04 [-0.12; 0.05]	-0.32 [-0.99; 0.35]
$r = 6 : k = 5$	-201.01*** [-251.87; -150.13]	0.12** [0.05; 0.18]	0.12* [0.03; 0.20]	1.26*** [0.57; 1.95]
AIC	23433.15	-1408.74	-1355.66	
BIC	23511.84	-1330.05	-1276.97	
Log Likelihood	-11702.57	718.37	691.83	
Deviance	23405.15	-1436.74	-1383.66	
Num. obs.	2040	2040	2040	
Num. groups: turk id	109	109	109	
Num. groups: experiment id	85	85	85	
Variance: turk id.(Intercept)	1381.80	0.03	0.03	
Variance: experiment id.(Intercept)	1535.68	0.00	0.00	
Variance: Residual	4953.09	0.02	0.02	

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table VII. Coefficients of Models Fit to Social Networking Game: Payoffs 2

	Payoff	Coop. Likelihood	Coop. Likelihood (Max Experience)	Coop. Likelihood (by round)
(Intercept)	372.44*** [291.21; 453.24]	0.80*** [0.73; 0.87]	0.79*** [0.72; 0.87]	3.3*** [1.68; 4.92]
experience	0.61 [-0.81; 2.06]	0.00 [0.00; 0.00]		0.19*** [0.14; 0.25]
max experience			0.00 [0.00; 0.00]	
round				-0.3*** [-0.16; -0.45]
experience:round				-0.019*** [-0.014; -0.024]
$k = 5$	302.38*** [214.58; 390.20]	0.09*** [0.07; 0.12]	0.09*** [0.07; 0.12]	1.55*** [1.12; 1.98]
$k = 23$	375.07*** [292.71; 457.31]	0.07*** [0.05; 0.10]	0.07*** [0.05; 0.10]	1.19*** [0.79; 1.59]
AIC	2792.20	-461.42	-461.49	
BIC	2820.04	-433.57	-433.65	
Log Likelihood	-1388.10	238.71	238.75	
Deviance	2776.20	-477.42	-477.49	
Nub. obs.	240	240	240	
Num. groups: turk id	53	53	53	
Num. groups: max experience1	35	35	35	
Num. groups: experiment id	10	10	10	
Variance: turk id.(Intercept)	0.00	0.00	0.00	
Variance: max experience1.(Intercept)	7071.50	0.01	0.01	
Variance: experiment id.(Intercept)	3151.99	0.00	0.00	
Variance: Residual	4599.19	0.00	0.00	

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table VIII. Coefficients of Model Fit to Wildcat Wells Data

	Payoff	Prob. Exploring	Prob. Exploring (Max Experience)	Prob. Exploring (by round)
(Intercept)	512.33*** [482.84; 542.00]	0.72*** [0.68; 0.76]	0.72*** [0.67; 0.78]	2.63*** [2.31; 2.95]
experience	-0.08 [-0.31; 0.15]	0.00 [0.00; 0.00]		0.014*** [0.011; 0.016]
max experience			0.00 [0.00; 0.00]	
round				-0.17*** [-0.16; -0.18]
experience:round				-0.0011*** [-0.001; -0.0013]
max. max. betweenness	-29.41*** [-46.01; -12.80]	0.00 [-0.02; 0.03]	0.00 [-0.02; 0.03]	-0.07 [-0.15; 0.01]
min. avg. betweenness	-50.24*** [-66.82; -33.64]	0.03** [0.01; 0.05]	0.03** [0.01; 0.05]	0.14*** [0.06; 0.23]
max. max. closeness	-34.33*** [-50.86; -17.79]	0.02* [0.00; 0.04]	0.02* [0.00; 0.04]	0.07 [-0.01; 0.15]
min. max. closeness	-27.62** [-44.24; -10.96]	0.01 [-0.01; 0.03]	0.01 [-0.01; 0.03]	0.12** [0.04; 0.2]
max. avg. clustering	-23.66** [-40.27; -7.03]	0.01 [-0.01; 0.03]	0.01 [-0.01; 0.03]	0.1* [0.02; 0.19]
min. avg. clustering	-57.72*** [-74.30; -41.10]	0.02 [-0.01; 0.04]	0.02 [-0.01; 0.04]	0.02 [-0.06; 0.1]
max. var. constraint	-51.55*** [-68.12; -34.95]	0.02 [0.00; 0.04]	0.02 [0.00; 0.04]	0.07 [-0.02; 0.15]
AIC	78077.33	-1327.04	-1327.87	
BIC	78164.37	-1240.00	-1240.83	
Log Likelihood	-39025.66	676.52	676.93	
Deviance	78051.33	-1353.04	-1353.87	
Num. obs.	5975	5975	5975	
Num. groups: user id	130	130	130	
Num. groups: experiment id	51	51	51	
Num. groups: trial	8	8	8	
Variance: user id.(Intercept)	2485.13	0.03	0.03	
Variance: experiment id.(Intercept)	6306.03	0.00	0.00	
Variance: trial.(Intercept)	260.98	0.00	0.00	
Variance: Residual	26213.38	0.04	0.04	

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$